

104

THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

Y 4. SCI 2:104/32

The High Performance Computing and...

HEARING

BEFORE THE

SUBCOMMITTEE ON BASIC RESEARCH

OF THE

COMMITTEE ON SCIENCE

U.S. HOUSE OF REPRESENTATIVES

ONE HUNDRED FOURTH CONGRESS

FIRST SESSION

OCTOBER 31, 1995

[No. 32]

Printed for the use of the Committee on Science



U.S. GOVERNMENT PRINTING OFFICE

21-158CC

WASHINGTON : 1995

For sale by the U.S. Government Printing Office
Superintendent of Documents, Congressional Sales Office, Washington, DC 20402

ISBN 0-16-052412-1

THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

Y 4. SCI 2:104/32

The High Performance Computing and...

HEARING

BEFORE THE

SUBCOMMITTEE ON BASIC RESEARCH

OF THE

COMMITTEE ON SCIENCE

U.S. HOUSE OF REPRESENTATIVES

ONE HUNDRED FOURTH CONGRESS

FIRST SESSION

OCTOBER 31, 1995

[No. 32]

Printed for the use of the Committee on Science



U.S. GOVERNMENT PRINTING OFFICE

21-158CC

WASHINGTON : 1995

For sale by the U.S. Government Printing Office
Superintendent of Documents, Congressional Sales Office, Washington, DC 20402
ISBN 0-16-052412-1

COMMITTEE ON SCIENCE

ROBERT S. WALKER, Pennsylvania, *Chairman*

F. JAMES SENSENBRENNER, JR., Wisconsin	GEORGE E. BROWN, JR., California RMM*
SHERWOOD L. BOEHLERT, New York	RALPH M. HALL, Texas
HARRIS W. FAWELL, Illinois	JAMES A. TRAFICANT, JR., Ohio
CONSTANCE A. MORELLA, Maryland	JAMES A. HAYES, Louisiana
CURT WELDON, Pennsylvania	JOHN S. TANNER, Tennessee
DANA ROHRABACHER, California	PETE GEREN, Texas
STEVEN H. SCHIFF, New Mexico	TIM ROEMER, Indiana
JOE BARTON, Texas	ROBERT E. (Bud) CRAMER, JR., Alabama
KEN CALVERT, California	JAMES A. BARCIA, Michigan
BILL BAKER, California	PAUL McHALE, Pennsylvania
ROSCOE G. BARTLETT, Maryland	JANE HARMAN, California
VERNON J. EHLDERS, Michigan**	EDDIE BERNICE JOHNSON, Texas
ZACH WAMP, Tennessee	DAVID MINGE, Minnesota
DAVE WELDON, Florida	JOHN W. OLVER, Massachusetts
LINDSEY O. GRAHAM, South Carolina	ALCEE L. HASTINGS, Florida
MATT SALMON, Arizona	LYNN N. RIVERS, Michigan
THOMAS M. DAVIS, Virginia	KAREN McCARTHY, Missouri
STEVE STOCKMAN, Texas	MIKE WARD, Kentucky
GIL GUTKNECHT, Minnesota	ZOE LOFGREN, California
ANDREA H. SEASTRAND, California	LLOYD DOGGETT, Texas
TODD TIAHRT, Kansas	MICHAEL F. DOYLE, Pennsylvania
STEVE LARGENT, Oklahoma	SHEILA JACKSON LEE, Texas
VAN HILLEARY, Tennessee	WILLIAM P. LUTHER, Minnesota
BARBARA CUBIN, Wyoming	
MARK ADAM FOLEY, Florida	
SUE MYRICK, North Carolina	

DAVID D. CLEMENT, *Chief of Staff and Chief Counsel*

BARRY BERINGER, *General Counsel*

TISH SCHWARTZ, *Chief Clerk and Administrator*

ROBERT E. PALMER, *Democratic Staff Director*

SUBCOMMITTEE ON BASIC RESEARCH

STEVEN SCHIFF, New Mexico, *Chairman*

JOE BARTON, Texas	PETE GEREN, Texas
BILL BAKER, California	ALCEE L. HASTINGS, Florida
VERNON J. Ehlers, Michigan	LYNN N. RIVERS, Michigan
GIL GUTKNECHT, Minnesota	LLOYD DOGGETT, Texas
CONSTANCE A. MORELLA, Maryland	WILLIAM P. LUTHER, Minnesota
CURT WELDON, Pennsylvania	JOHN W. OLVER, Massachusetts
ROSCOE G. BARTLETT, Maryland	ZOE LOFGREN, California
ZACH WAMP, Tennessee	MICHAEL F. DOYLE, Pennsylvania
DAVE WELDON, Florida	SHEILA JACKSON LEE, Texas
LINDSEY O. GRAHAM, South Carolina	(Vacancy)
VAN HILLEARY, Tennessee	(Vacancy)
SUE MYRICK, North Carolina	

*Ranking Minority Member

**Vice Chairman

CONTENTS

WITNESSES

	Page
October 31, 1995:	
Dr. Anita K. Jones, Director, Defense Research and Engineering, and Chair, Committee on Information and Communications, National Science and Technology Council	6
John Toole, Director, National Coordination Office for High Performance and Computing Communications	21
Dr. Ivan E. Sutherland, Vice President and Fellow, Sun Microsystems, and Co-chair of the National Research Council's Committee to Study High Performance Computing and Communications	34
Dr. John Ingram, Research Fellow, Schlumberger, Austin, Texas	214
Dr. Edward D. Lazowska, Professor of Computer Science and Engineering, University of Washington, and Board Member, Computing Re- search Association, Seattle, Washington	223
Dr. Jeremiah P. Ostriker, Charles A. Young Professor of Astronomy, Provost, Princeton University, Princeton, New Jersey	275
Dr. Forest Baskett, Chief Technology Officer and Senior Vice President of Research and Development, Silicon Graphics, Inc., Representing the Computer System Policy Project	281

THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

TUESDAY, OCTOBER 31, 1995

U.S. HOUSE OF REPRESENTATIVES,
COMMITTEE ON SCIENCE,
SUBCOMMITTEE ON BASIC RESEARCH,
Washington, DC.

The Subcommittee met, pursuant to notice, at 9:33 a.m. in room 2318, Rayburn House Office Building, Hon. Steven H. Schiff [Chairman of the Subcommittee] presiding.

Mr. SCHIFF. Good morning. I would like to convene this hearing of the Basic Research Subcommittee and invite the members to join me in their places, and I would like to begin with a brief opening statement.

I would like to welcome everyone to this hearing on the High Performance Computing and Communications Program.

When this program was started in 1991, very few of us had heard the terms we hear so often today. the Internet, world wide web, global information infrastructure, information society, and cyberspace.

Today we have supercomputers using models for climate prediction, the design of supersonic aircraft, and to recreate the surface of distant planets, as well as the design and understanding of nuclear weapons which drove the supercomputer's initial development.

Today every presidential candidate has his own home page posted on the Internet.

With the current speed of technology advancement, a generation of technology can be obsolete in the same amount of time it takes Congress to introduce, debate, and make a decision on a bill.

As we all know in this new Congress, we are faced with tough funding decisions. The American public has sent us a clear message. They want reduced spending and less government. However, they also expect us to find the right roads of investment that will continue to make America secure and support the research that will provide the economic drivers of the future.

Unfortunately, all too often in past Congresses, programs that were developed for short periods of duration have ended up with eternal life.

With the HPCC program spending over \$1 billion this year, I believe the cost must be justified in the current budget climate.

With that in mind, the focus of this hearing is on three questions.

What is the current status of the HPCC Program?

Second, is there a need to reauthorize the HPCC Program next year when the act expires? Alternately, the act could be authorized on an annual basis.

Number three, if we authorize the program, what recommended changes should be made to the program?

Today we have two panels testifying before us. The first panel's witnesses will brief the Subcommittee on the current program, actions of the Federal Government in HPCC, and present outside evaluation of the program.

The second panel of witnesses will discuss industry's use of HPCC, and we will hear from scientists involved in the program and their recommendations for the future direction of HPCC.

Before I recognize our first panel of witnesses, I would like to welcome and recognize the Subcommittee's ranking minority member, Mr. Geren of Texas, for any opening statement he might like to make. Mr. Geren.

Mr. GEREN. Thank you very much, Mr. Chairman.

Characterization of the emerging post-industrial world as the Information Age has now become commonplace. The significance of the High Performance Computing and Communications Program, which we are reviewing this morning, is that the program is focused on research areas that will advance the technologies which are driving this transformation.

The idea for the HPCC Program originated with the research community itself in the late 1980's when it had become apparent to researchers that significant advances in computing power and electronic communications were possible through closer collaboration among related Federal R&D programs. The Science Committee encouraged these efforts to create a focused, multi-agency research initiative to develop more powerful computing and network capabilities and to develop the applications software needed to exploit the capabilities of advanced computers and networks for solving challenging problems in science and engineering research.

Authorizing legislation developed by the Science Committee in the House was subsequently enacted with broad bipartisan support in 1991. President Bush's budget request for fiscal year 1992, in anticipation of approval of the authorizing legislation, proposed a 30 percent increase for the computing and networking R&D budgets of the Federal agencies which had been planning the HPCC Program. Subsequently, both the Bush and Clinton administrations proposed funding for the program consistent with the authorization levels in the 1991 act.

The Science Committee held several oversight hearings during the past Congress to assess the progress of the HPCC Program and to consider expanding the scope of the program which has now been done, to include research and demonstration activities on applications of HPCC technologies in areas of broad public benefit, such as K through 12 education, access to government information, and health care delivery. Overall, nearly uniform support for the program is expressed in testimony received by the Committee from the computer, telecommunications, and information industries, from manufacturing and service industries which rely on high performance computing technology, and from all parts of the academic research and education communities.

To obtain an in-depth evaluation of the program, Congress turned to the National Academy of Sciences to conduct a study that was concluded earlier this year. The results of the study, which will be presented to the subcommittee this morning, identifies specific accomplishments of the program, including making parallel computing a practical approach for achieving high performance and developing user-interface software for easily finding and accessing information via the Internet. The National Academy also has recommended that the program be continued.

This morning we will once again hear from a range of Government and private sector witnesses who have been asked to assess the accomplishments of the program and to advise the Subcommittee on the need for its reauthorization. We are interested in exploring whether the program remains focused on the most important research issues and whether it ought to continue in its current form. Recommendations for improvements to the current authorizing statute are especially sought.

Mr. CHAIRMAN. This program is one of the most significant Federal research initiatives, and I am pleased that you have called this oversight hearing. I would like to join you in welcoming our witnesses this morning and I look forward to their testimony. Thank you.

Mr. SCHIFF. Thank you, Mr. Geren. I just want to reiterate it is always a pleasure to work with you and your staff in this Subcommittee.

Mr. GEREN. Thank you, Mr. Chairman.

Mr. SCHIFF. We are joined by our full Committee ranking Democratic member, Congressman Brown of California. Congressman Brown, do you desire to make an opening statement?

Mr. BROWN. Mr. Chairman, I have an opening statement, but in deference to the importance of the subject and our witnesses, other than congratulating you on this very important hearing, I will put the rest of it in the record.

[The prepared statement of Mr. Brown follows:]

PREPARED STATEMENT OF HON. GEORGE E. BROWN, JR.

I congratulate the Chairman for calling this oversight hearing on the High Performance Computing and Communications Program. During a year in which general disparagement of Federal programs has been much in evidence, it is refreshing to contemplate this federally sponsored research program on information technology, because this is a field that will contribute in a major way to the future prosperity of the Nation and because it has been driven forward by the successes of Federal R&D over the past 50 years.

The origins of digital computers and computer networks are closely tied to federally funded R&D. Moreover, many of the advances in computing and networking that have been exploited by American companies—or that provided the foundation for start-up companies that became major corporations—sprang from federally sponsored basic research and exploratory development activities. Some of these connections between federally funded research discoveries and specific commercial developments are highlighted in the recent National Academy of Sciences report on the HPCC program, which we will hear more about this morning. An interesting observation about these connections is that the path from basic research result to application is not linear and the time from basic result to application averages from 10 to 15 years.

This observation underlines the importance of a Federal role in R&D, even for a healthy, rapidly developing technology area, such as high performance computing and communications. That is, to succeed when new product cycles may be only 18 months long, a company will tend to focus its R&D on the next few product cycles. It will not support much longer term, generic research because such basic research

is broad in applicability—the connection to an eventual application may be obscure—and a particular research finding may require many iterations to perfect as a marketable product. Consequently, the research results inevitably become public so that the company cannot recoup its research investment.

The HPCC program was designed to address long-term, generic research questions. While industry's needs are taken into account in determining broad program goals, the research is driven by the participating scientists and engineers at academic and government laboratories, where 80 percent of the research is performed.

The evidence thus far, including the findings of the National Academy's review of the program and past hearings before the Science Committee, is that the High Performance Computing and Communication Program has been largely successful. I look forward to the opportunity this morning to hear more about the program's accomplishments and to obtain the views of our witnesses on how the program may be strengthened.

Mr. SCHIFF. Thank you, Mr. Brown.

I want to advise all members, without objection, any written opening statement they might want to offer will be made part of this hearing, but I do not want to foreclose brief opening remarks from any member of the Subcommittee who might want to make them. So, let me look first over to my side of the aisle. I see no requests. Let me see on the other side. Mr. Doyle?

Mr. DOYLE. Thank you, Mr. Chairman. I do have an opening statement, but in the interest of time I will submit that statement for the record. Thank you very much.

Mr. SCHIFF. Again, all written statements will be made part of the record.

[The prepared statements of Mr. Doyle and Mrs. Morella follow:]

PREPARED STATEMENT OF HON. MIKE DOYLE

I want to thank Chairman Schiff for holding today's hearing on the High Performance Computing and Communications Initiative and welcome today's distinguished witnesses.

As HPCC nears the end of its initial 5-year authorization, this hearing provides a timely opportunity to examine what it is that has been accomplished and how we should proceed from here with HPCC. I welcome the recent National Research Council report and believe that it will be a useful tool for this Committee to use in its deliberations. I do have some reservations about the report which I expect will be discussed in today's hearing. However, I strongly endorse the NRC's conclusion that,

"...strong public support for a broadly based research program in information technology is vital to maintaining U.S. leadership in information technology. Facilitated access for both academic and industrial users to advanced computing and communications technologies has produced further benefits both in scientific progress and in U.S. industrial competitiveness."

While the HPCC has enjoyed strong bipartisan support throughout the life of the program, I am concerned that there are those who, despite the progress our Nation has achieved through this program, will ignore the facts and seek to dismantle it. Already, in this Congress, HPCC has been targeted—a \$35 million reduction for NASA, one of the major contributors to HPCC, was included in the House version of the VA-HUD Appropriations bill. This Committee, too, acted to cut NASA's HPCC budget during the consideration of H.R. 2043, without the benefit of hearings. I submitted Additional Views to the Report accompanying H.R. 2043 which I would like to submit for the record of this hearing as well.

I am hopeful that Members will take the time to understand how this program works, as well as how it has positively impacted both the American economy and our academic institutions. Furthermore, in assessing the worth of this program, I urge my colleagues to pay special attention to the positive impact HPCC has had on the cost of government operations. For example, due to the advancements which have resulted from HPCC, the Department of Defense is able to develop imaging systems that operate at much greater cost-effectiveness; also, the National Security Agency makes great use of supercomputing in its codebreaking operations.

While there is need for some finetuning of HPCC, as pointed out in the National Research Council report, it would be a decisive step backwards for Congress to

abandon it either out of ignorance or as a result of election-year trophy hunting. Unfortunately, I think the best course of action for us to pursue, a stand-alone authorization bill, would be one which the course of action which leaves HPCC most susceptible to these types of attacks. I am disappointed that the climate towards a large number of our Federal science programs in the 104th Congress seems to suggest that we are incapable of rationally reviewing vital R&D programs.

Despite what I see as serious setbacks for science and technology in our country over the past few months, I remain hopeful that there continues to be sufficient bipartisan support for HPCC. Thus, I hold out hope that we could move an authorization bill expeditiously through this Committee and through the House without it being subject to questionable characterizations of being an example of "corporate welfare"—a term that has come to mean many things to many people. Let us examine HPCC intelligently and honestly without turning one of our most important and successful programs into a political football.

ADDITIONAL VIEWS SUBMITTED BY HON. MIKE DOYLE ON NASA'S HIGH PERFORMANCE COMPUTING & COMMUNICATIONS PROGRAM

I am greatly concerned about the adverse impact on computer and communications research that will result from the Committee's action to cut \$35 million from the NASA High Performance Computing and Communications (HPCC) Program.

Federal funding under HPCC primarily supports university-based research, which underpins the development of new computer and communications technologies for use in engineering, financial services, manufacturing, medicine, security, space, and other areas. This fundamental research usually lies outside the research and development funded by the communications industry, as it involves long-term investment and uncertain financial payoffs. Cuts to the HPCC program reduce the scale of basic research in critical technology areas and ultimately will cripple the ability of the U.S. to compete internationally in computer and communications technologies.

Since its inception under the High Performance Computing Act of 1991, the HPCC Program has received bipartisan support and has performed well under Committee review. The legislative report (H.Rept. 102-66, Part 1) states that, at the full Committee mark-up of the High Performance Computing Act on May 8, 1991, "a bipartisan amendment in the nature of a substitute, sponsored by Representatives Boucher, Valentine, Brown, Packard, Lewis, and Walker, was adopted."

In an October 26, 1993, hearing the Committee received testimony from the manufacturing, financial services, energy and aerospace industries supporting the value of the program. An oversight hearing on May 10, 1994, revealed no serious problems with HPCC. The most recent assessment of the program was carried out by the National Academy of Sciences, which praised HPCC in a report earlier this year stating:

"Accomplishments under the High Performance Computing and Communications Initiatives to date reveal two key trends: better computing and computational infrastructure and increasing researcher-developer-user synergy. In the committee's expert judgement, the program has been generally successful."

The HPCC program is closely coordinated among 12 Federal R&D agencies with many jointly funded activities. NASA provides approximately 12 percent of the total HPCC budget and is the agency responsible for the coordination of the government, academic, and industrial partners in the advanced software technologies component of the initiative.

The Committee has held no hearings on this program during the 104th Congress. New Members of Congress, who comprise a majority of the Committee on Science, have not had the opportunity to learn about the objectives and accomplishments of this multi-agency, closely coordinated program. In the absence of any hearing record in this Congress, major cuts in the HPCC program are, at best, ill-advised. A large cut, such as the one contained in this bill, will have a negative impact on the ability of the overall HPCC program to meet its goals.

As a program that has proven its success, HPCC merits continued Federal funding at a level that will allow it to continue its activities. The decision to drastically reduce the NASA component of the HPCC budget should be reconsidered.

PREPARED STATEMENT OF HON. CONSTANCE A. MORELLA

Mr. Chairman, I wish to commend you for your leadership in this important issue.

As you know, I share your interest with the High Performance Computing and Communications (HPCC) program, and its peripheral development of the National Information Infrastructure and I look forward to hearing the testimony of our witnesses today.

Just this past summer, this Subcommittee and my Technology Subcommittee, jointly reviewed one of the latest related developments on this issue—pornography on the Internet. The hearing was very useful and informative in reviewing the technologies currently available to assist parents in restricting children's access to pornographic materials on the Internet. I am pleased to have explored that timely issue with you.

Now, we move on to another timely topic—the possible reauthorization of the HPCC Act. This hearing will provide us with the status of the HPCC program and recommendations for its future direction. I am interested in the testimony of today's panel and in their views regarding the efficacy of the current HPCC structure; the possible effect of the absence of a direct authorization upon the HPCC program; whether the Federal-private partnerships are being leveraged effectively with Federal funds or whether the program could proceed without Federal funds; which of the 12 Federal agencies participating in the program are adequately doing their job and which ones are not; and specific overall recommendations to improve the management and coordination of the program.

Thank you, Mr. Chairman.

Mr. SCHIFF. Before turning now to the witnesses, I have to make one observation. I suspect that certainly the people who are the witnesses are accustomed to this, but for anyone in the room who does not have happened to have seen a congressional hearing before, there will be some comings and goings by the members. I have to offer a couple of amendments in the Judiciary Committee sometime this morning on a bill that is being marked up, and other members have similar other commitments because a number of things go on at the same time.

I just want to stress that the importance of any hearing is the record that is being made by the court reporter, and that in due course all Members of the House of Representatives will have made available to them the entire contents of this hearing, and that is what is most important about a hearing.

With that having been said, I would like to welcome our first panel of witnesses. They are Dr. Anita Jones who is Director of Defense Research and Engineering and Chair of the Committee on Information and Communications of the National Science and Technology Council; Mr. John Toole who is Director of the National Coordination Office for High Performance Computing and Communications; and Dr. Ivan Sutherland who is Vice President and Fellow of Sun Microsystems and Co-chair of the National Research Council's Committee to Study High Performance Computing and Communications.

We want to welcome you all here and I want to say that, without objection, all of your written statements, if submitted, will be made part of the record. I invite you to proceed with your oral testimony in summary in any way that you wish. Dr. Jones, I would like to begin with you, please.

STATEMENT OF DR. ANITA K. JONES, DIRECTOR, DEFENSE RESEARCH AND ENGINEERING, AND CHAIR, COMMITTEE ON INFORMATION AND COMMUNICATIONS, NATIONAL SCIENCE AND TECHNOLOGY COUNCIL

Dr. JONES. Mr. Chairman, members of the Subcommittee, and staff, I am very pleased to be here to speak to you about the Federal High Performance Computing and Communications Program.

I am here in the capacity as Chair of the NSTC Committee on Information and Communications, as you said, but I also bring the perspective of the Department of Defense.

I would like to address the issues you raised about the importance of, the evolution of, and the future of the high performance computing program and am particularly pleased to be here with Mr. Toole, Dr. Sutherland, and with some additional members of the agencies who are major players in the high performance computing program.

This program is the centerpiece of the Federal Government's information technology research and development investment. High performance computing and communications are discriminators. They permit the construction of high resolution models of both artificial worlds and physical reality which in turn offer a way, an alternative to theory, to experimentation for scientists and technologists to pursue knowledge and understanding. High performance computing and communications enable scientists and engineers to solve problems that could not be solved yesterday.

Every several years the question, what is high performance computing, and the question, what is high performance communications, have different answers, at least measured in terms of speed. At any given time the highest performance computers and networks are in the research lab where they are expensive to develop and most difficult to apply. But one generation's highest performance computing will be mastered in the laboratory and will be harnessed for the highest priority needs of the Nation in the next generation. And then in the generation beyond, they will be ubiquitously and cost-effectively applied in the laboratory, in the office, and in some cases in the home. This has happened again and again over the past several decades and it reflects the vitality of American science and American industry.

The generation of research must come first. If the United States is to maintain the lead in high performance information technology that it has today, that early investment must be made. Some say, let industry do it. In my view that is not a logical request. Industry's objective is to return a profit, not to invent new technology. Industry will seek low risk paths to the goal of profit in order to protect the interest of shareholders and to maximize the probability of success. Where a suite of marketed products already exist, a corporation will seek incremental product enhancement to maintain and expand market share. Companies that define new ventures select technology that already appears ripe and ready to lead to a return on investment. In the United States, funding the search for solutions to the very long-term, the very far-future problems falls to the Government, and the information technology research investment of the Federal Government has provided a manifold return over the years.

Advancement in science and technology is a competition. It can be won but it does not stay won. It is a continual competition and certainly we in the Department of Defense understand that. Research results prepare the way for the next scientific advance, the next agency mission advance, and the next decade's industrial dominance. There is a causal link between the federally funded research and information technology through the High Performance

Computing and Communications Program and the Nation's scientific prowess and its industrial prowess in information technology.

This vitality has but it requires a Government and industry partnership where Government funds the long-term, high-risk research and industry aggressively exploits new opportunities. In the past, the U.S. has dominated high performance computation, but that competition has increased over the years. Today the top 20 highest performance computers in the world are half in the United States and half in Japan. It is a competition that is a competition at the research level, not just at the industrial level.

The High Performance Computing and Communications Program arose out of an ongoing federally funded research program across many areas and many agencies. Congress sounded a clarion call for coordination on a national scale in this arena. It has worked. It continues to be highly effective today. The High Performance Computing and Communications Program is the very best implementation of the notion of a virtual agency that I have encountered. The coordination which now involves 12 participating agencies is no longer an initiative, but a routine way of aggressively pursuing the business of those agencies. Agency representatives who come together in the National Coordination Office communicate, they set agendas, they leverage the funding and results across each other's programs, and at the same time, I see no dilution of agency mission or agency autonomy. As the proverb goes, "The whole is greater than the sum of the parts."

The multiple agency cooperation has enabled the creation of projects that would not otherwise be possible. The gigabit testbed effort is just one example. When that was started, industry was not performing that kind of organized research activity and Government was an important catalyst to the experiment that ensued. The Government often galvanizes the new research arena, such activity, such experimental activity, that accelerates progress and research and eventually leads to acceleration in industrial harnessing of technology.

The content of the High Performance Computing and Communications Program has evolved over time. Early effort was on very high speed hardware, on very high speed networks, and on the grand science challenges, selected fundamental problems in science and engineering that have broad both scientific and economic impact. That program's focus has evolved over time. Earlier we extended it to include the national challenges, selected, fundamental applications that have broad and direct impact on the Nation's competitiveness and the well-being of its citizens. Today the program continues to effectively and aggressively evolve.

The National Science and Technology Council's Committee on Information and Communications, which has some oversight responsibility of this program, has developed a strategic plan that outlines the next succession of broad areas of future investment. This process involved academia, industry, and Government to assess the best ideas of all and we have defined six strategic focus areas. They include things like high performance/scalable systems which continues to be a major area of activity in high confidence systems.

These focus areas, though, span all performance ranges. However, the high performance contexts are the most challenging because of the speeds and the capabilities involved and because of the problems you can address at the high end that you cannot address at the modest and low end. There continues to be a succession of new problems to be solved at the physical hardware level, in packaging, systems software, applications software, at system interface levels, and in communications protocols. Large scale multimedia data storage, more capable human-system interaction, and predictable scalability are all important challenges.

In your invitation to testify here today, you asked me to cite examples of how this program supports the mission responsibilities of the participating agencies. I included in my testimony a set of eight varied and different examples. This program has had material impact on essentially all the participating agencies in the program. The Federal investment has helped to establish the systems, the centers, the pools of human expertise, the mathematics solutions, the algorithms for particular applications, the visualization tools to solve the agency mission applications.

Let me highlight a couple.

One is the Department of Energy science-based nuclear stockpile stewardship. International treaties now forbid operational tests of nuclear weapons. Yet those weapons remain our ultimate deterrent. The Nation is presented with a unique challenge; How to ensure that the current weapon systems remain functional, how to upgrade them, how to engineer replacements in the absence of physical testing. The results of the last several decades of research make it possible to even consider using numeric simulation as a solution to those challenges. It has made possible this approach.

A second example is weather prediction, of interest to certainly the Department of Defense, but a mission of NOAA, NSF, and NASA. Recent devastations caused by multiple hurricanes and storms in the Southeast underscore the necessity of greatly improved weather forecasting. The prediction of Hurricane Opal was less than we could have wanted. Significant improvements over today's state of the art will be possible only with higher capacity computers and better algorithms to model the weather. Weather prediction has a necessary real-time delivery requirement. You have got to complete the simulation before the weather happens. According to the National Weather Service, predicting a disaster in time can save as much as \$1 million per mile of evacuation costs. Accurate, timely predictions can prevent needless evacuations and allow for more orderly ones.

I grew up in the Gulf Coast. We never listened to the Weather Service requests for evacuation. That is no longer true. Citizens evacuate when the Weather Service says bad weather is coming because they are almost always right.

Medicine is another example. The NIH seeks solutions to merging the new sensor images that we can now take of a body in hopes of learning new information out of the merged images that is greater than the information we can gain out of separate images. This, too, requires new achievements at the results level.

Let me close discussing contributions of the high performance computing program by highlighting that one unique result that it

returns to the Nation is the education of scientists and engineers. The program educates scientists and engineers. New curricula have been brought to place computational science in the classrooms broadly. But the High Performance Computing and Communications Program has produced students in computer architecture, parallel programming, language design and implementation, mathematical support for parallel computation, visualization, and management of networks. Those individuals, that expertise is necessary to solve the mission agency applications and to provide industry with technical leaders that they need in the future.

The High Performance Computing and Communications Program can be credited with many accomplishments. The interagency coordination has been effective. The agencies have a strong commitment to investment in research and technology development in this area, particularly at the high end as an integral part of their agency program. Consequently, we seek authorization not of a continued initiative but a continuation of the core program of the individual agencies who strongly seek this committee's authorization support in the form of individual agency program element authorizations. And we will continue to work closely with your staffs in any way necessary to ensure this bipartisan investment can continue.

I believe that it is crucial to mission agencies that we continue the research investment. Several weeks ago President Clinton relaxed the export control regulation so that higher speed computers could be freely exported out of the United States than in the past. However, the very highest speed computation engines manufactured by U.S. industry may not be exported in order to protect our national security. It is because high performance computation, when applied to agency mission problems and our national industry problems, can give us unique advantage.

To highlight that, I would cite the performance in Desert Storm of the F-117, the stealth fighter. It flew only 5 percent of the missions. It is responsible for 30 percent of the target kills. It is planar in structure. It is a set of planar surfaces in its outward appearance. One reason for that is at the time of its design we did not have the algorithms, but, in particular the high performance computing that enabled us to do a continuous surface design. You see in the B-2 a continuous surface, a more effective stealth aircraft. We had better high performance computing capabilities because of this program.

Our objectives both in stealth signature for future aircraft and in other areas of military application require better mathematics, better algorithms, higher performance computers, and solutions to the problems being broached by the continued High Performance Computing and Communications Program.

Federal investment, steady, sustained, and well-managed, is necessary to maintain the U.S. leadership in high performance information technology as it progresses at a very rapid pace.

Mr. Chairman, Subcommittee members, I appreciate the opportunity to discuss this critical program with you. We strongly seek your support. We need that support. I would be pleased to address any questions that you might have.

[The prepared statement of Dr. Jones follows:]

Prepared Statement by :

Dr. Anita K. Jones

Director, Defense Research And Engineering

**Chairman,
Committee on Information and Communications,
National Science and Technology Council**



to the

SUBCOMMITTEE ON BASIC RESEARCH

HOUSE COMMITTEE ON SCIENCE

October 31, 1995

**Not for Publication until
Released by the House Committee
on Science**

DEPARTMENT OF DEFENSE
PREPARED STATEMENT BY

DR. ANITA K. JONES
DIRECTOR OF DEFENSE RESEARCH AND ENGINEERING
TO THE
SUBCOMMITTEE ON BASIC RESEARCH
OF THE
HOUSE COMMITTEE ON SCIENCE
OCTOBER 31, 1995

ON

FEDERAL HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS (HPCC)
PROGRAM

Mr. Chairman, Members of the Subcommittee, and staff: I am pleased to be here today to speak to you about the Federal High Performance Computing and Communications (HPCC) Program. I am here in the capacity of chair of the National Science and Technology Council (NSTC) Committee on Information and Communications. Additionally, I serve as the Director of Defense Research and Engineering in the Department of Defense. I will address the issues that you have raised about the importance, evolution, and future of the HPCC Program. I am particularly pleased to appear with Dr. Ivan Sutherland who co-chaired the recent National Research Council study on HPCC requested by Congress and with Mr. John Toole who last spring became the first full-time director of the National Coordination Office for HPCC.

The High Performance Computing and Communications Program is the centerpiece of the Federal Government's information technology R&D investments. Today, information and communications technology are a major contributor to the nation's capability in science, health care, education, national security, and a major factor in our nation's economic strength. But it is high performance computing and communications that is a discriminator. They permit the construction of high-resolution models of both artificial worlds and physical reality, which in turn offer an alternative to theory and experimentation in the pursuit of knowledge and understanding.

High-performance computing and communications enables scientists and engineers to solve problems that could not be solved yesterday.

Every several years the definition of what is "high performance computing" and "high performance communications" substantially increases in terms of speed. At any given time the highest performance computers and communications are in the research laboratory where they are expensive to develop and most difficult to apply. But one generation's highest performance computing will be mastered in the laboratory, and will be harnessed for the highest priority needs of the nation in the next generation. And then in the generation beyond, they will be ubiquitously and cost-effectively applied in the office and possibly the home. This has happened again and again over the past several decades.

The generation of research must come first. If the United States is to maintain the lead in high performance information technology, the early investment must be made. Some say "let industry do it". In my view, that is not a logical request. Industry's objective is to return a profit, not to invent new technology. Industry will seek low risk paths to that goal in order to protect the interest of their shareholders and to maximize the probability of success. Where a suite of marketed products already exist, a corporation will seek incremental product enhancement to maintain and expand market share. Companies that define new ventures will select technology that already appears ripe, and ready to lead to a return on investment. In the United States funding the search for solutions to the very long-term, far-future problems falls to the government. And the information technology research investment of the federal government has provided a manifold return.

The HPCC Program goals include:

- * Extension of U.S. leadership in high performance computing and networking technologies;
- * Dissemination of these technologies to accelerate innovation and serve the economy, national security, education and the environment; and,
- * Enhancement of gains in U.S. industrial competitiveness.

The Program has, and continues, to serve those goals, and to achieve them. But competition in technology is not "won"; there is a continual competition. Research results prepare the way for the next scientific advance, the next agency mission advance, and the next decade's industrial dominance. There is a causal link between federally funded research and the Nation's industrial prowess in information technology. It was stated in the recent review of the High

Performance Computing and Communications Program by the National Research Council, entitled "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure." Government funded research in tools and techniques for the application of parallel systems facilitated their use to solve agency mission problems and eventual adoption by industry. Similarly government support for the training of graduate students made expertise in high performance computing available for hire by industry.

In rapidly advancing technologies, government and industry are in partnership. Government funds the long term, the high risk research, and industry aggressively exploits the new opportunities to build market and to take global leadership. If anything, the long term research competition has become more competitive. In the past, the U.S. dominated high performance computation. Today, of the top twenty highest performance computers in the world, half are in the U.S. and half are in Japan. If government investment in high performance computing research falters then I believe we will have reduced opportunity for U.S. industry to sustain what continues today to be our lead.

The HPCC Program arose out of an on-going federally funded research program across many areas. Congress sounded a clarion call for coordination on a national scale. It worked. It continues to be highly effective today. The HPCC Program is the best implementation of the notion of a "virtual agency" that I have encountered. The coordination, which now involves 12 participating agencies, is no longer an "initiative", but routine business. Agency representatives communicate, set agendas, leverage the funding and the results of each other's program. At the same time, I see no dilution of agency mission or agency autonomy. As the proverb says "the whole is greater than the sum of the parts."

This multiple agency coordination enables the creation of projects that would not otherwise be possible: for example, the gigabit testbeds are conducted on a scale that would be impossible for one agency to field. Transition of technology between agencies occurs more readily and rapidly because of the coordination and the familiarity of one agency on an on-going basis with the programs and expected results of another agency's programs. As I said, coordination of high performance communications and computing is now part of the way our agencies conduct normal business.

The content of the HPCC Program has evolved over time. Early focus was on teraflop (very high speed parallel) computers, high-speed networks (one billion bits per second) and the Grand Challenges, which are selected, fundamental problems in science and engineering that have

broad scientific and economic impact. Later the program's focus was enlarged to encompass the National Challenges, which are selected, fundamental applications that have broad and direct impact on the Nation's competitiveness and the well-being of its citizens.

Enabling applications have been both important to solve, and important in that they exhibit the software challenges that need to be overcome in order to apply high performance hardware resources broadly. Software tools, techniques, and algorithms are an important focus today as is the "scaling" of software solutions so that they run cost effectively on a range of sizes of systems. The applications are as varied as the agency missions involved in the HPCC Program. And the Nation's HPCC Program continues to evolve.

The NSTC Committee on Information and Communications (CIC), which I chair, has developed a Strategic Implementation Plan that outlines broad areas of future investments. We have opened this planning process up by engaging academia, industry, and government to assess our mutual ideas and approaches. These inputs are being used to formulate our strategy and investments for the future.

To ensure that research and underlying technology development is fully responsive to end-user applications and to national goals, activities must be strategically focused and efficiently coordinated. Together with each of the Agencies, the CIC has developed six Strategic Focus Areas to guide federal research and technology investment in information and communications into the next century. Strategic focus areas represent key opportunities to focus, coordinate, and accelerate information and communications science and technology development.

The Strategic Focus Areas are:

- * Global-Scale Information Infrastructure Technologies
- * High Performance/Scaleable Systems
- * High Confidence Systems
- * Virtual Environments
- * User-Centered Interfaces and Tools
- * Human Resources and Education

These focus areas span all performance ranges. However, the high performance contexts are most challenging because of the speeds and the capacities involved. There continues to be a succession of new problems to be solved at the physical hardware level, in packaging, systems software, applications software, intra- and inter-system interfaces, and in communications

protocols. Large-scale, multimedia data storage, more capable human - system interaction and predictable scalability are challenges.

These are just some of the problems which must be solved to meet the NSTC CIC goal to "accelerate the evolution of existing technology and nurture innovation that will enable universal, accessible, and affordable application of information technology to enable America's economic and national security in the 21st century." The evolution signifies both success at past challenges and a vital program.

Research to Support Agency Missions

In your invitation to testify here today, you asked me to cite examples of how the HPCC Program supports the mission responsibilities of the participating Federal Agencies. Essentially, the HPCC Program has blazed a trail that facilitates mission application. Within the HPCC Program advanced computing centers were created whose resources are used to solve versions of the Grand and National Challenges. The software tools and techniques that are developed by the various researchers, centers, and consortia in the HPCC Program collectively make high performance computers and networks more accessible and usable as a tool to solve mission problems.

The following list provides examples of selected agency mission applications. Each one has and will rely upon the result of the HPCC Program to achieve its ends.

I. (DOE) Science-based Stockpile Stewardship

International treaties now forbid operational tests of nuclear weapons. Yet nuclear weapons remain our ultimate deterrent. The Nation is presented with a unique challenge: how to ensure that current weapon systems remain functional, and how to upgrade and engineer replacement systems in the absence of physical testing. This challenge is being addressed by the Department of Energy's Accelerated Strategic Computing Initiative (ASCI). Projections of computing requirements to perform the needed modeling and computational experimentation exceed by more than a thousand times the capacity of the world's largest computer.

ASCI will extend the DOE's ability to test and to prototype nuclear weapon systems using numerical simulation. To meet the nuclear stockpile stewardship needs of the year 2010, today's two-dimensional simulations of weapon responses in certain special case conditions must be

extended to achieve higher-resolution, three-dimensional simulations of all physical phenomena and all systems' behavior. Advances in simulation requires both much higher performance computers as well as software tools and techniques to harness the computation resource together with better application codes.

2. (NOAA, NSF and NASA) Weather Prediction

Recent devastation caused by hurricanes Andrew, Hugo, Opal, and other tropical storms in the Southeast underscore the necessity of greatly improved weather forecasting. The predictions of hurricane Opal were particularly poor in identifying the landfall point and the speed of progress. Significant improvements over today's state of the art will only be possible with even higher capacity computers than exist today. Weather prediction has an unusual real-time delivery requirement: simulations must run faster than real time to be timely. According to the National Weather Service, predicting a disaster in time can save as much as \$1M per mile of evacuation costs. Accurate, timely predictions can prevent needless evacuations, and allow for more orderly ones when necessary.

3. (NASA) Aeronautics

Computer simulations for commercial and space vehicles are credited with reducing in half both design time and cost, as well as improving performance. Additional savings were achieved through a 1.5 percent reduction in Specific Fuel Consumption for a large commercial aircraft engine; a fleet of aircraft using that engine realizes millions of dollars in fuel cost savings over the fleet's lifetime.

Looking to the future, NASA experimental and operational vehicles will be performing in transonic and hypersonic regimes where current state-of-the-art computer simulations are unable to accurately simulate complicated turbulent effects today. Fully accurate simulations must be multi-disciplinary, incorporating not only fluid dynamics computations but also structural, thermal, and chemical effects. Such computations, on full-scale aircraft geometry's, are well beyond the capability of currently available supercomputer systems.

4. (DoD) Stealthy Fighters and Bombers

The F117 stealth fighter that performed so effectively in Desert Storm - it flew 5% of the sorties and killed 30% of the designated targets - is planar in shape. One reason for this is that at

the time of its design, smooth continuous surfaces could not be accurately modeled with the computer and software simulations at hand. The B2, a more modern stealth aircraft, has smooth surfaces. Improvements in high performance computing was one enabler of this advance.

The development of future stealth vehicles relies crucially on the capability of computing radar cross-sections of vehicle designs in the multi-GHz frequency range. Today, the state of the art is limited to small vehicle sections. To perform similar computations on full-scale aircraft will require higher performing computers together with more accurate modeling of the physical phenomena.

Other emerging defense applications of high-performance computing and communications include modeling and simulation for training purposes, battlefield medical technology, battle tactics simulation, rapid response planning and logistics, and design and cost optimization of weapon systems.

5. (NSA) Cryptology and Signal Processing

The exploitation of methods for the analysis of encrypted messages can consume all envisioned computation capabilities. Sensor derived signal data is growing in volumes at a rapid rate not just for intelligence applications but for medical, geophysical and myriad other applications. The processing of such data share with weather predication the need for timely processing. Higher performance alone will not solve the demand. New software techniques to focus attention on the region with value are essential.

6. (NSF) Computational Biology

The community of biologists is deciphering the patterns of deoxyribo nucleic acid (DNA) material that guide development of the individual human being. This is called the human genome project. It requires high performance computing to determine how material whose elements are known, fold to form a three dimensional structure. This is call the protein folding problem and the computer is used to analyze the multitude of candidate structures. Understanding this three-dimensional structure is need in studying the function of proteins in living systems and designing new ones for biological and medical purposes.

7. (NASA and NSF) Astrophysics

One of the outstanding open problems of astrophysics and cosmology is to model the evolution of the universe from the big bang to its present-day form. To date, simulations have had only moderate success. Astrophysicists are uncertain whether this is due to fundamental errors in our physical understanding of the universe, or to inadequacies in the numeric simulations. More accurate simulations await more powerful computer systems, algorithm and software advances.

8. (NIH) Medicine

Medical imaging requires high performance computing to manipulate and analyze images. Such images are routinely used today. One promising opportunity involves merging the data from multiple image sensors (e.g., positron emission tomography (PET), magnetic resonance imaging (MRI), computerized axial tomography (CAT)). This requires high performance computers and new software techniques. It is hoped that the merged representation of the body will yield more useful information, when analyzed, than any single image.

There is one additional contribution that the HPCC Program has made to the agencies with mission challenges and to industry with market competitive challenges. And that is the education of scientists and engineers. The HPCC Program continues to educate scientists and engineers in the context of research in high performance computing and communications. New curricula have brought computational science into classrooms at all levels and increased the diversity of students pursuing scientific disciplines. The HPCC Program has produced students trained in areas such as computer architecture, parallel programming languages, mathematical support for parallel computation, scientific visualization, and the design and management of very high speed networks. In addition, students are trained in the computational science aspects of many scientific application areas, a continuing pipeline of new human expertise is vital to continued national leadership.

Recommendation and Summary

High performance computing and communications is a subject for scientific study itself. At the same time is an enabler for federal agency missions and our national industrial competitiveness. The global competition in scientific research and in the development of technology in this area is a continuing competition. Sustained investments in information and communications are critical to success.

The HPCC Program can be credited with many accomplishments. The interagency coordination is effective. The agencies have a strong commitment to investment in research and technology development in the information technologies, particularly at the high end, as part of their agency program. Consequently, we seek authorization not of a continued initiative but of the core program. We strongly seek Authorization Committee support in the form of individual Agency program element authorizations. We will continue to work closely with your staffs in any way necessary to ensure this bipartisan investment can continue.

Several weeks ago President Clinton relaxed the export control regulation so that higher speed computers could be freely exported from the U.S. than in the past. However, the highest speed computation engines manufactured by U.S. industry may not be exported in order to protect our national security. High performance computation can solve problems effectively and rapidly that cannot be solved by a nation that lacks that computational capability. To capitalize on the advantage held by the U.S. requires continued scientific discovery and development. For example achieving militarily superior stealth signatures for future aircraft will rest in part on higher performance computation and better software than is available today. And, there are comparable opportunities in science and in industry.

Federal investment, - steady, sustained and well managed - is necessary to maintain leadership in high performance information technology as it progresses at a rapid pace. I see ample evidence that our leadership in science, our national security, and our economic health are all strengthened by such federal investment. There is no substitute.

Mister Chairman and Subcommittee Members, I appreciate the opportunity to discuss this critical program with you. I would be pleased to address any questions you might have.

Mr. SCHIFF. Thank you, Dr. Jones.
Mr. O'Toole. Mr. Toole. Excuse me.

STATEMENT OF JOHN TOOLE, DIRECTOR, NATIONAL COORDINATION OFFICE FOR HIGH PERFORMANCE AND COMPUTING COMMUNICATIONS

Mr. TOOLE. Thank you, Mr. Chairman. A lot of people think I am Irish, so it is no problem at all.

[Laughter.]

Mr. TOOLE. I am really happy to be here before the Chairman, members of the Subcommittee, and staff to tell you about the accomplishments, plans, and directions of the Federal HPCC Program. After serving 9 years at ARPA, the Advanced Research Projects Agency, I am now the first full-time Director of the National Coordination Office for High Performance Computing and Communications. I believe we have an unprecedented opportunity for establishing the foundation for America's information future and would like to talk to you about that today.

Information technology, as Dr. Jones highlighted, is really central to our national security and to our society. Today's initial national information infrastructure, global information infrastructure, and things that we read and hear on TV and elsewhere are really based on technologies emerging from the computing and communications research that has been underway for many years. However, realizing the dreams of tomorrow will really rest upon the investments we make in research today.

The Federal HPCC Program has been a model virtual agency since its inception and really has maintained an important balance between the very long-term computer science research and the most advanced computing and information technologies possible, all directed at the collective mission needs of the respective agencies.

The National Coordination Office, which I now direct, was created to coordinate the Federal HPCC Program across the Federal Government. In fiscal year 1996, 12 Federal departments and agencies will participate in the HPCC Program by coordinating their R&D activities and accelerating technology transfer into key, computationally intensive, and information intensive application areas. The estimated 1995 HPCC Program budget for the 12 agencies was \$1,084 million. In fiscal year 1996, the President requested \$1,143 million for the 12 organizations.

Projects are competitively selected, and overall approximately 52 percent of the program monies go to academia, 20 percent to industry, and 28 percent to Government labs, not-for-profit institutions.

In the short time I have today, I would like to focus on three basic topics. First, some examples of key program accomplishments that highlight the new science national capabilities and future infrastructure; secondly, a brief synopsis of critical R&D investments that I think must be made, particularly those that are best served by interagency efforts; and third, recommendation to the Committee on reauthorization, as you requested.

I am really pleased to provide the Committee copies of the report, High Performance Computing and Communications. Foundation for America's Information Future. It is also on-line and I think

we provided each of the members and staffers copies of this. It really documents in detail the significant program accomplishments commonly called the Blue Book and outlines the breadth of current activities. Although we have produced these reports every year of the program to highlight our accomplishments, this year produces and shows many on-line activities that people can get even additional information for the reader that may be interested.

You probably have heard of interdisciplinary Grand Challenge R&D projects. These are computationally intensive applications that have led the new science and engineering techniques in a wide range of disciplines. Modeling air flow and turbulence around aircraft, properties of the engines, combustion, the oceans, the atmosphere, structural dynamics of car crashes, and the evolution of galaxies, and innovation in high performance computing and communications techniques have brought the Nation new knowledge and new capabilities.

The spectacular images transmitted around the world of the recent collision of the Shoemaker-Levy 9 comet fragments with the planet Jupiter illustrate the impact of HPCC technologies across multiple Federal agencies. It is an example of new scientific computational models run on a new generation of scalable computing systems communicated almost instantaneously on a web of networks that made possible the use of NASA's Hubble space telescope in unforeseen ways. In essence, it allowed us to predict what was really happening and to observe this phenomenal event.

In the future we can expect that scientific results, developed as part of Grand Challenges, combined with what we have called National Challenges, which are information intensive applications impacting U.S. competitiveness and societal well-being, will greatly enhance the quality of life for Americans. Weather forecasting and medical information are two good examples. Dr. Jones cited the weather example.

In the medical field, applications of telemedicine are really bringing the physician instantly in contact with patients in remote locations using advanced computer networks. Digital models of human anatomy are being developed to provide a new education tool for researchers, health care providers, students, and the general public.

Networking has been an important component of the HPCC Program and will really help determine the future of where we are going in information technology. The success of the Internet we have heard and read is widely known beginning with ARPA-funded research in the 1960's and continuing to NSF's privatization of the Internet today. A major activity of this program has been six gigabit testbeds jointly funded by several HPCC agencies to demonstrate the use of high performance communications technology for agency missions. Nine Federal agencies, 13 telecommunications carriers, 12 universities, eight corporations, and two State computer centers have participated in these six testbeds that have connected 24 sites.

Asynchronous transfer mode, ATM, and synchronous optical networking, SONET, technologies, developed rapidly in these collaborative projects, and prototype switches and protocols were de-

veloped to address issues at the very high speed that probably would have taken many years longer without the HPCC Program.

High performance computing systems have been essential to national defense, NASA, NSF, NIH, and the Department of Energy scientists. Research has been ongoing on high performance backplanes, operating systems, embedded systems, networks of workstations, and new algorithms. In addition, we of course recognize the great speed requirements and computational speed such as Sandia's 281 gigaflop linear algebra benchmarks performed on the Intel Paragon computers.

Software has always lagged hardware development. It is challenging to write efficient codes for complex computations on new machine architectures. However, portable languages, such as high performance FORTRAN, high performance C++, along with computational and software tools to support them, are available for use and emerging from the research funded out of this program.

The HPCC Program added a fifth component in 1994, information infrastructure technology and applications. This has supported enabling information infrastructure research and development based on related research in high performance computing and communications. One of the most notable achievements were the web browsers used to retrieve vast amounts of information available around the world. As a leading prototype, the mosaic system developed at the National Center for Supercomputing Applications really became part of a model that became used and has now been copied many, many millions of times by many different companies and is being used as an enabler across the whole network.

Applications as part of the HPCC refer to experimental use of advanced information technology research and development applied to real problems in innovative ways. These applications drive the HPCC technology development, and by engaging in Grand and National Challenges in the HPCC Program, agencies have a unique opportunity to accelerate the incorporation of these technologies into their mainstream mission activities. For example, a 4-year digital library research program being jointly supported by HPCC agencies will lead to new ways to store, retrieve, search, and process massive amounts of information for the future.

Educational resources and computational tools from the HPCC Program are also important results of this scientific investment. Students, including K through 12, undergraduate, graduate, and post-graduate students, are now approaching the scientific and engineering tasks with computational tools that were never available before. Researchers and students, for example, are capable of studying the human genome structure, biological processes at the cellular and molecular levels, and evolution of galaxies, and the design of new materials. I think you will hear about some of this in the second panel as well.

As Dr. Jones had pointed out and I would like to reiterate, the Federal program has had broad impact on many other Government programs who have become users of these important technologies from science-based stockpile stewardship to defense systems modernization. DOE's accelerated strategic computing initiative is an important example of a critical mission initiative that is based upon the long-term research of the HPCC Program. Simply stated,

a vigorous R&D initiative focused on the high end is essential to their success.

I would like to move to some of the future directions and to highlight and continue where Dr. Jones hinted at in terms of the strategic focus areas.

We had developed a strategic implementation plan, again which we have also provided each of the Committee members, that sets broad, general areas of investment for the future. We have used this strategic implementation plan as a guide and have engaged academia, industry, and Government in really a revolutionary way and for us to help define our future R&D investments that will lead us into the Information Age. The most notable forum was in July and results have been available on the Internet, and we have an ongoing dialogue with these different communities.

In addition, each agency has been intensely involved in its own program and budget planning which will ultimately lead to the President's budget for fiscal year 1997.

CIC has also moved to meet the interagency management and organizational challenges.

First, the National Coordination Office, which I direct, is broadening its role to support the CIC.

Secondly, CIC is in the process of creating an applications council which will move cutting edge technologies into mission critical applications across Government and to provide a feedback mechanism channel between CIC and the Federal user community.

Third, the Federal Networking Council has been placed under CIC to provide even better networking collaborations across Federal agencies and closer synergy with the R&D managers.

Lastly, the NCO is physically moving to a more central location, thanks to the gracious generosity of NSF that will become our host, and we expect to move there in about a month.

Now I would like to highlight some of the key long-term R&D areas that are essential to realize our information future.

First, there is a need for cost-effective high performance computing. High performance/scalable parallel processing is now part of the strategic plans of many Federal agencies and U.S. industries. However, there are numerous barriers to satisfying computing needs beyond the foreseeable future. Scientific and technical computing is at risk unless those barriers are overcome and long-term R&D is needed to overcome those barriers. Today many in the computing systems industry are looking to use commodity off-the-shelf components in building computing systems in the future. Continued research aimed at innovation and continued innovation in component technologies and systems, for example, will be important for us at the turn of the century. As part of these high end efforts, the HPCC program has established a task force to detail the needs and options for such investments.

Second, increased emphasis is being placed on critical software research for scalable parallel computing systems. Today more emphasis is being put on the software tools and development environments. We have established a second task force in this area to develop and coordinate the multi-agency investments in this area.

Third, networking research will emphasize high performance protocols, services embedded in the network, and other ways to

achieve an open data network as described in the National Research Council report *Realizing the Information Future—the Internet and Beyond*. We are stimulating the development of a truly inter-operable systems environment for the future that all America's agencies can participate in.

Fourth, it is essential that we build high confidence systems that provide information security, privacy, reliability, and invulnerability. This must include all aspects of the system, including hardware communications networks and software. An interagency group is looking at some of this again as part of the HPCC Program and people beyond the scope within the CIC.

Fifth, the program will support demonstrations of advanced applications that combine Grand and National Challenges and use the underlying technology base that we are investing in. Through the use of HPCC technologies, new discoveries in science and engineering are emerging. We are also demonstrating new information infrastructure technologies and applications. Key enabling technologies, such as advanced digital libraries, which I cited before, and information management, will provide important insight into the future.

You asked us to address the reauthorization issue. The importance of information technology to everyone I think here has been completely evident, but the real revolution is yet to come. The HPCC Act of 1991 was very important legislation. It established a multi-agency initiative focused on a critical national issue. It also demonstrated a strong working partnership in this area between Congress and the administration.

Although another major reauthorization bill could be useful, we do not feel it is necessary in order to continue this important program. Over the past 4 years, each of the agencies has become committed to information technology and the interagency process to ensure coordination across them. I believe we have responded to the challenge. We believe that continued strong congressional support of each agency's HPCC activities and the HPCC coordination processes will ensure the most efficient program execution for the next 5 years.

Mr. Chairman, I am acutely aware of the budget pressures facing the Committee and the need to reduce our Federal deficit sharply. I also believe, however, that information technology is our future and Government investments we are making in computing and communications will shape our long-term ability to succeed. I firmly believe that our investments, coupled with unleashing the brilliant skills of academia and industry, are the keys to our future for our children and our grandchildren.

I thank you for the opportunity to address the Committee. I also have with us representatives from each of the four agencies that you described, Dr. Howard Frank from ARPA, Dr. Paul Young from NSF, Mr. John Cavallini from the Department of Energy, and Mr. Lee Holcomb from NASA, for later questions if you would like.

[The prepared statement of Mr. Toole follows:]

Prepared Statement by:

JOHN C. TOOLE

Director, National Coordination Office for High Performance Computing and Communications

**Chairman, High Performance Computing, Communications, and Information Technology Subcommittee,
Committee on Information and Communications,
National Science and Technology Council**

Before the

SUBCOMMITTEE ON BASIC RESEARCH

HOUSE COMMITTEE ON SCIENCE

October 31, 1995

**Not for Publication until
Released by the House Committee
on Science**

Federal High Performance Computing and Communications (HPCC) Program

October 31, 1995

Introduction

Good morning Mr. Chairman, Members of the Subcommittee, and staff. I am pleased to appear before this subcommittee and have the opportunity to describe the accomplishments, plans, and directions of the Federal High Performance Computing and Communications (HPCC) Program. I am very honored to have taken responsibility as Director of the National Coordination Office (NCO) for HPCC in March 1995 from Dr. Don Lindberg, who served concurrently as Director of NCO and as Director of the National Library of Medicine from September 1992 until March 1995. After serving the past 9 years at the Advanced Research Projects Agency, I am now the first full-time Director of the NCO. I believe we have an unprecedented opportunity for establishing the foundation for America's information future and would like to talk to you about that today.

Information technology is central to our national security and to our society, both economically and socially. US leadership has resulted from an extraordinarily complex and fruitful long-term partnership among academia, industry, and government. Today's initial National and Global Information Infrastructures are based on technologies emerging from the computing and communications research that has been underway for many years; however, realizing the dreams of tomorrow will rest upon the investments we make in research today.

The Federal HPCC Program has been a model "virtual agency" since its inception, with unpreceded collaboration among the (currently) 12 Federal organizations that participate in HPCC R&D. The Program has maintained an important balance between the very long term computer science research, and the most advanced computing and information technologies possible -- all directed at the collective mission needs of the respective agencies. Agencies have strongly supported the program because it is a complex field in which advancements are in their strategic best interests, and collaboration provides the most effective leverage possible.

The Program began during the Bush Administration and received strong bipartisan support in both the House and Senate. Since it began, the Program has been an Administration priority for our long term science and technology investment, providing long term capability for the next wave of information technology while exploiting high performance computing and communications across the globe. The National Coordination Office, which I now direct, was created to coordinate the Federal HPCC Program across the Federal government.

In FY 1996, 12 Federal departments and agencies will participate in the HPCC Program by coordinating their R&D activities and accelerating technology transfer into key computationally intensive and information-intensive application areas. The estimated FY 1995 HPCC Program budget for the nine participating Federal organizations was \$1,084 M. For FY 1996, the President requested \$1,143 M for the 12 organizations. Projects are competitively selected and overall, approximately 52% of the program moneys went to academia, 20% to industry and 28% to government labs/not for profits.

I am honored to appear today with Dr. Anita Jones and Dr. Ivan Sutherland. Dr. Sutherland was one of the co-chairs of the recent study by the National Research Council, *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*. This was a major study requested by Congress. In addition to looking at the HPCC Program in detail, the NRC committee of distinguished researchers from academia and industry studied the long term investments made by government and industry. This committee concluded that the government's investments in information technology have made a significant impact over the long term. The HPCC Program is the most critical part of that investment today.

In the short time I have today, I'd like to focus on three topics:

- Examples of key Program accomplishments that highlight new science, national capabilities, and future infrastructure.
- A brief synopsis of the major critical R&D investments that must be made, particularly those that are best served by interagency efforts.
- Recommendation to the committee on reauthorization, as you requested.

Program Accomplishments

I am pleased to provide the committee copies of the report, *High Performance Computing and Communications: Foundation for America's Information Future*, which can also be viewed on-line via the World Wide Web. This report, prepared by the HPCCIT Subcommittee of the National Science and Technology Council's Committee on Information and Communications R&D (CIC), describes in detail the Program's significant accomplishments and outlines the breadth of current activities. Although we have produced such reports every year, this year's report includes on-line links to many of the highlighted research projects, providing even more detail for the reader. I'll use a couple of examples to highlight some of the key investments that have been made since the Program's formal authorization in FY 1992.

Interdisciplinary Grand Challenge R&D projects -- computationally-intensive applications - have led to new science and engineering techniques in a wide range of disciplines. By modeling air flow and turbulence around aircraft, properties of their engines, combustion, the oceans, the atmosphere, the weather, pollution, climate, groundwater, earthquakes, vegetation, the human body, proteins, enzymes, the human brain, materials, chemicals, structural dynamics of car crashes, and the evolution of galaxies, innovation in high performance computing and communications techniques have brought the Nation new knowledge and new capabilities. The HPCC program has set some leading edge examples for many scientific disciplines to follow.

Research in computational modeling is starting to pay off, for example, in the US aeronautics industry. A comprehensive new engine modeling system called the Numerical Propulsion Simulation System, was developed under NASA using these techniques. It has resulted in engineering productivity improvements that enabled one of our premiere aircraft engine companies to cut design time in half for high-pressure jet engine compressors that are used in the Boeing 777. This new design also reduces fuel consumption -- saving billions of dollars in fuel cost over the life of the fleet and reducing harmful environmental impact.

The spectacular images transmitted around the world of the recent collision of the Shoemaker-Levy 9 comet fragments with the planet Jupiter illustrate the impact of HPCC technologies across multiple Federal agencies. It is an example of enabling new scientific

computational models, run on a new generation of scalable computing systems, communicated almost instantaneously on a web of networks, that made possible the use of the NASA's Hubble space telescope in unforeseen ways. NSF's Pittsburgh Supercomputing Center and DOE's Sandia National Laboratories accurately foretold the event using new computational techniques that depend very much on high performance computing technology. These models provided the information for space scientists to point the Hubble telescope in the correct direction to observe first hand these unique events. Without the modeling efforts and the underlying computing research, a once-in-a-lifetime observational astronomy event would have been lost.

In the future, we can expect that scientific results developed as part of the Grand Challenges, combined with National Challenges -- information-intensive applications impacting US competitiveness and societal well being -- will greatly enhance the quality of life for Americans. Weather forecasting and medical information are two good examples. Using modern techniques requiring high performance computing, NOAA has developed a hurricane prediction system that can more accurately predict the path of a hurricane, provide earlier warning, and, in turn, save lives. In the medical field, application of telemedicine is bringing the physician instantly in contact with patients in remote locations using advanced computer networks. Digital models of human anatomy are being developed to provide a new education tool for researchers, health care providers, students, and the general public.

Internetworking has been an important component of the HPCC Program, and will help determine the future of the NII. The success of the Internet is widely known, beginning with ARPA-funded research in the 1960's and continuing to NSF's privatization of the Internet today. For the research community and individuals across the Nation, both at work and at home, the Internet ties people and information resources together to work on future challenges. HPCC research in networking is focused on the future capabilities, such as technical approaches for performance, scale, and security.

A major activity has been the six gigabit testbeds jointly funded by several HPCC agencies to demonstrate the uses of high performance communications technology for agency missions. Nine Federal agencies, thirteen telecommunications carriers, twelve universities, eight corporations, and two state supercomputer centers have participated in these six testbeds that connect 24 sites. Asynchronous Transfer Mode (ATM) and Synchronous Optical Networking (SONET) technologies, developed rapidly in these collaborative projects, and prototype switches and protocols were developed to address issues at very high speed that would have taken many years longer without the HPCC Program. Research has led to opportunities for industry to open unforseen markets and be at the forefront of technology in the 21st century.

During the past year, the HPCC program accomplished an important milestone in interagency networking. The Department of Energy and NASA have worked together with a US telecommunications partner to implement the first fast packet Internet. This illustrates how HPCC agency partners continue to successfully collaborate, make new capabilities available to their communities, and are transitioning away from a dedicated infrastructure.

High performance computing systems have been essential to national defense, NASA, NSF, NIH, and Department of Energy scientists. The HPCC program has funded research to explore new systems approaches and scalable techniques for a wide class of problems. Research has been on-going on high performance backplanes, operating systems, embedded systems, networks of workstations, and new algorithms. In addition, we acknowledge the impressive world records in computation speeds, such as Sandia's 281 gigaflops linear algebra benchmarks performed on the Intel Paragon computers.

Software development has always lagged hardware development. It is challenging to write efficient software for complex computations on new machine architectures; however, portable languages such as High Performance FORTRAN and High Performance C++, along with software tools to support them, are available for use in new computational experiments.

Experimental versions of these tools are now available to the public over the National HPCC Software Exchange, an activity supported by the HPCC Program to ensure rapid dissemination of advanced software to researchers throughout the US. Associated with the software development activities are new visualization techniques, such as the Cave Automatic Virtual Environment, which allows a user to explore new design approaches and unique ways to visualize the massive amounts of data today's researcher has available.

The HPCC program added a fifth component in FY 1994, Information Infrastructure Technology and Applications. In the past two years, this component has supported enabling information infrastructure R&D, based upon related research in high performance computing and communications. One of the most notable achievements are the "Web browsers" used to retrieve vast amounts of information available all around the world. Based on the Mosaic system developed at the National Center for Supercomputing Applications, funded as part of NSF's HPCC program, already more than 1 million copies of public domain Mosaic software have been obtained from NCSA; more than 10 million copies of Enhanced NCSA Mosaic have been licensed through commercial start-up companies such as Spyglass, and there are many others such as Netscape, Spry, etc. This R&D activity at NCSA started an unforeseen segment of the communications industry, and spawned many commercial companies that will make global information accessible to all citizens and enterprises. The technologies and experience emerging from the program have been instrumental in formulating a long term view of the future National Information Infrastructure and how the US might evolve in the information age.

Applications as part of HPCC refer to the experimental use of advanced information technology research and development applied to real problems in innovative ways. These applications drive the HPCC technology development. By engaging in Grand and National Challenges in the HPCC Program, agencies have a unique opportunity to accelerate the incorporation of these critical technologies into their mainstream mission activities. For example, the four-year digital library research being jointly supported by HPCC agencies will lead to new ways to store, retrieve, search, and process masses of information being gathered by weather satellites, census surveys, and many other types of relevant information. Health care, national security, energy management, aeronautical design, public health, and education all benefit from the technologies of the HPCC program.

Educational resources and computational tools from the HPCC Program are an important result of this scientific investment. Students, including K-12, undergraduate, graduate, and post-graduate students, are now approaching these scientific and engineering tasks with computational tools that were never before available. Researchers and students, for example, are now capable of studying the human genome structure, biological processes at the cellular and molecular levels, the evolution of galaxies, and the design of new materials, by using computational modeling. Projects supporting the classroom, such as NASA's Classroom of the Future, provide a new generation of educational modules and teacher support, using remote sensing databases over the Internet. More and more people are coming to the conclusion that research began under the HPCC Program is fostering a revolution in the way we do science, the way we learn, and the way we share information.

Even though some of the innovative computer manufacturers in the dynamic and highly competitive market of high performance computing and communications may fail, their insight and research into the nature of parallel computational algorithms and the structure of future

computational systems have survived. Indeed, they brought new knowledge and insight into the field of high performance computational science and engineering.

Finally, the HPCC Program has helped support a world-class network of research facilities that are reaching out to the states, communities, and individuals. I call your attention to the descriptions of the High Performance Computing Research Centers described in the FY 1996 Supplement to the President's budget that I have provided for you today. These Centers are actively engaged in research, education, outreach to minority institutions and rural areas, and training our next generation of students throughout the US.

As Dr. Jones will cover in her testimony, the Federal HPCC Program has had broad impact on many other major government programs who have become users of these important technologies - from science-based stockpile stewardship to Defense Systems modernization. DOE's Accelerated Strategic Computing Initiative (ASCI) is an important example of a critical mission initiative that is based upon the long term research of the HPCC program. Simply stated, a vigorous R&D initiative focused on the high end is essential to their success. Such examples clearly show that information technology is one of the key enablers for virtually all mission-driven applications.

Future Directions

As part of the National Science and Technology Council (NSTC), the Committee on Information and Communications R&D (CIC) was established under the leadership of the Honorable Anita Jones, Director of Defense Research and Engineering and CIC chair, and the Honorable Lionel S. Johns, Office of Science and Technology Policy, CIC co-chair. The CIC, which consists of senior representatives from Federal R&D Agencies, has published and put on line a strategic implementation plan titled *Committee on Information and Communications Strategic Implementation Plan: America in the Age of Information*. This plan outlines several strategic focus areas designed to focus fundamental information and communications research and to accelerate development in ways that are responsive to NSTC's overarching goals, agency mission goals, and our Nation's long term economic and defense needs. The strategic focus areas are: global-scale information infrastructure technologies, high performance/scalable systems, high confidence systems, virtual environments, user-centered interfaces and tools, and human resources and education.

Using this strategic implementation plan as a guide, we have engaged academia, industry, and government in different fora to help define R&D investments that will lead us into the information future. The most notable event was a forum held in July - results of which have been available on the Internet for public comment and feedback. In addition, each agency has been intensely involved in its own program and budget planning, which will ultimately be reflected in the President's budget for FY97.

CIC has also moved to meet the interagency management and organizational challenges. First, the National Coordination Office (NCO) is broadening its role to support CIC. Second, CIC is in the process of creating an Applications Council. This Council will provide a means to move cutting-edge technologies into mission critical applications across government and to provide a feedback channel between the CIC and the Federal user community. Third, the Federal Networking Council (FNC) has been placed under the CIC to provide even better networking collaborations across Federal agencies and closer synergy with the R&D managers. Fourth, the NCO is physically moving to a more central location. NSF has graciously agreed to become our host, and we expect to move in about a month.

Now I'd like to highlight some of the key long term government R&D areas that are essential to realize the information future. Sustained investment over the long term is essential, particularly with the inherent shorter term focus of industry. It is hard to predict which new ideas and approaches will succeed, since the exact course of exploratory research cannot be planned in advance, and progress in the short term is difficult to quantify. Furthermore, the industrial R&D investment, although much larger than government investment, is very different in nature, necessarily focusing primarily on shorter term product development cycles. The rationale for government investments in this crucial area are developed in detail in the NRC HPCC study.

US strength in information technology - both its scientific and economic, competitive strength - is due in critical ways to the aggressive federal research investment in computing, communications, and agency mission applications. While enormous progress is now visible in high performance computing and communications in the first four years of the Program, much remains to be done. These needs have led to several new directions being taken in FY 1996 - directions that build on these years of efforts:

First, there is a critical need for cost-effective high performance computing. High performance scalable parallel processing is now part of the strategic plans of many Federal agencies and US industries. However, there are numerous barriers to satisfying computing needs beyond the foreseeable future. Scientific and technical computing is at risk unless those barriers are overcome, and long term R&D is needed to overcome those barriers. US industry has a shorter-term focus with its R&D, but has done well absorbing research results produced by our strong university system supported by Federal research efforts. Today, many in the computing systems industry are looking to use "commodity off-the-shelf" components in building computing systems in the future. Continued research aimed at innovation in component technologies and systems will be important at the turn of the century. As part of these high end efforts, the HPCC Program has established a task force to detail the needs and options for such investments.

Second, increased emphasis is being placed on critical software research for scalable parallel computing systems including software tools, development environments, operating systems, languages, compilers, and programming libraries. Today, more emphasis is being placed on high performance software tools and development environments. This has arisen in part from open workshops in which academia, industry, and government have participated. We have established a second task force to develop and coordinate multi-agency investments in this area.

Third, networking research will emphasize high performance protocols (advances on today's Internet protocol), "services" embedded in the network, and ways to achieve an "open data network," as described in the 1994 National Research Council report, "Realizing the Information Future - the Internet and Beyond." This will stimulate the development of a truly interoperable systems that satisfy a wide range of user needs, and will provide greater diversity among communications media, such as "wireless" and interactive cable systems, for schools, libraries, health care facilities, and homes.

Fourth, it is essential that we build high confidence systems that provide information security, privacy, reliability, and invulnerability. This must include all aspects of the system, including hardware communications networks and software. An interagency HPCC group is formulating plans for this R&D activity.

Fifth, the Program will support demonstrations of advanced applications that combine Grand and National Challenges. Through the use of HPCC technologies, new discoveries in science and engineering are emerging. Both HPCC technologies and these new discoveries are

penetrating the science and engineering base of the Nation. The Program is also demonstrating new information infrastructure technologies and applications. For example, key enabling technologies such as advanced digital libraries and information management will provide important insight into the future. Accelerated research in areas such as user centered systems and applications to our Nation's most important needs, including the needs of the Federal government, can provide enormous leveraging to Federal funding of R&D while enhancing US economic competitiveness and quality of life in the 21st century.

Recommendation

The importance of information technology to everyone in the US has become evident, but the real information revolution is yet to come. The Program's accomplishments are substantial, but sustained investments by the Federal government are critical for future success.

The High Performance Computing Act of 1991 was important legislation - it established a multi-agency initiative focused on a critical National issue. It also demonstrated a strong working partnership in this area between the Congress and the Administration. Although another major reauthorization bill could be useful, we do not feel it is necessary to continue this important program. Over the past four years, each of the agencies has become committed to information technology and the interagency process to insure coordination across them. We believe that continued strong Congressional support for each agency's HPCC activities and the HPCC coordination processes will ensure the most efficient Program execution for the next five years.

Summary

Mr. Chairman, I am acutely aware of the budget issues facing this Committee and the need to reduce our Federal deficit sharply. I also believe, however, that information technology is our future, and government investments we make in computing and communications will help shape our long term ability to succeed. I firmly believe that our investments, coupled with unleashing the brilliant skills of academia and industry, are keys to the future for our children and grandchildren.

Thank you for this opportunity to address you. I welcome any questions or comments from the Committee.

Mr. SCHIFF. Thank you very much, Mr. Toole.
Mr. Sutherland, you are recognized.

STATEMENT OF DR. IVAN E. SUTHERLAND, VICE PRESIDENT AND FELLOW, SUN MICROSYSTEMS, AND CO-CHAIR OF THE NATIONAL RESEARCH COUNCIL'S COMMITTEE TO STUDY HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS

Dr. SUTHERLAND. Good morning, Mr. Chairman and members of the Committee. My name is Ivan Sutherland. I am Vice President and Fellow at Sun Microsystems which is a leading manufacturer of computer workstations.

In 1994, as you know, Congress asked the National Research Council to examine the status of the High Performance Computing and Communications Initiative. I helped respond to that request by serving as Co-chair with Fred Brooks of the resulting study committee that put together a report for your interests of the results of our study. That report was delivered early this year and I believe that the Committee has copies of it available today.

I am going to use two charts today, both of which come from that report.

As you probably know, the National Research Council is the operating arm of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The National Academy of Sciences was chartered by Congress in 1863 to advise the Government on matters of science and technology. I am a member of both the National Academy of Sciences and the National Academy of Engineering.

I would like to turn to my first chart here which is a graph, and what it shows is the increase in processing power, speed, per dollar over the last 5 years. Now, when an engineer sees a chart like this—it is plotted on a logarithmic scale. You can see that the scale goes from 10 to 100 to 200—he says, wow, that is amazing progress.

But you know what this chart really shows is just something that we have all observed. We all remember when fax did not exist and now cellular telephones are ubiquitous. Many of us remember when television was only black and white and now live video in color comes from around the world into our living rooms. We have all observed the explosive power of the Internet, and the power of yesterday's supercomputer is now in our homes, educating our children and entertaining all of us. For a few hundred dollars, you can buy a portable device that will tell you where you are in the world very accurately through the global positioning system with amazing precision, and it works anywhere on earth.

But you know what is truly amazing about the information machines is not how good they are but how quickly they get better. Every 15 years or so, they get 1,000 times better. That is what this chart basically shows.

Now, to understand what 1,000 times means, let us imagine riding across the plains in a covered wagon. It went about 20 miles a day. Then you have to imagine the early Model T Ford. It went about 20 miles an hour. Then you have to imagine the stealth fighter that Anita was talking about earlier which goes about 20

miles a minute. That is a factor of a thousand of progress. From 20 miles a day to 20 miles a minute is about a factor of 1,000.

Now, imagine that amount of progress in 15 years and that is what is amazing about the information business. The reason that is possible, of course, is that information is different from physical materials. Information does not weigh anything. Reproducing it does not have any inherent cost, and so it is possible to make this amazing progress in the information area.

Over the past 30 years, the progress in the information business has been driven by a revolution in computing in integrated circuits. Integrated circuits have gone from single transistors to chips with millions of transistors on them in remarkably few years. The rapid advance of integrated circuits is the fuel driving the information revolution. I anticipate that improved integrated circuits will continue to provide that fuel for at least another decade and probably longer. That means that we face between now and the early part of the next century probably another factor of 1,000 in the capability of our information machines.

That rapid pace of information technology presents a great opportunity. The nation that properly anticipates how best to use tomorrow's computing and communications machines gains great advantage in defense, commerce, medicine, education, and quality of life. Ever since World War II, we in the United States have enjoyed an advantage in computing and communications, producing new products, providing new services, forming new companies, creating new jobs to enhance our defense and to provide important export opportunities. Our ability to use the new devices quickly and well rests in part on our Yankee ingenuity, but it also rests in part on careful preparation that has been offered by an enlightened Federal program of support of long-term research in the information sciences.

Our HPCCI report contains a second chart, to which I would now like to turn. You have that. It is on page 2 and again on page 20 in that green report book.

This chart covers 30 years of progress in some of the aspects of information technology. On it we show nine of today's major information industry sectors, each of which grew out of Federal support of research in computing and communications. In each case we show the interplay between federal support in universities and Government laboratories and private industry support in industrial laboratories.

Now, if you look at the chart, you will see dark solid lines. That is the Federal support, and then you will see these lines that have diagonal stripes in them. That is the industrial support. You see arrows that go between various parts of the chart which are in fact showing the transfer of people or ideas from one of those research activities to the other. Then there are some bold, dotted lines, the dot-dot-dot-dot-dot part of the thing. That shows where that technology reached commercial product offerings. Then we go to racing stripes when that industry sector got to be a billion dollars per year of activity.

Here are nine areas, and you can see that they all have some roots in Federal funding. Some of them started in industrial things and then moved to a Federal funding situation. There is a very in-

teresting interplay between the industrial and the federally funded R&D here which I think has been the strength of this area.

Now, if you look at this chart, you can draw some lessons from it. In each case it took about 10 years or more for new ideas to reach market acceptability. It took an additional 5 years or so for the combined revenue of the markets to reach \$1 billion. In many cases brand new companies formed around the new ideas from federally supported research to bring the new products to market. As a Nation, of course, we get our return on the Federal research investment from the jobs, the taxes, and the exports produced by these companies, as well as from their products.

Now, I am pleased to work for one of the companies that appears on this chart. It is in the workstation area, and we are Sun Microsystems. We are a major manufacturer of computer workstations. Sun is only 13 years old and it has revenues of about \$6 billion per year. Sun was started by some young people who saw the future of computing and communications in federally supported research programs at Stanford and Berkeley.

Given the amazingly rapid pace of progress in information systems, how can we anticipate where information technology will lead 5 or 10 years or 15 years hence? By then machines with the power of today's supercomputers will appear in private homes. Perhaps such machines will understand speech. Perhaps they will recognize faces. Perhaps they will anticipate human needs far better than we can now imagine. We do not know exactly how they are going to help us or exactly what they are going to do, but we can be sure that the value of information machines 1,000 times more powerful than we have today will be enormous. They will do things we do not dream of today. Who in a covered wagon anticipated flying coast to coast in half a day or coming to Washington, as I did yesterday, and will return again to California this afternoon? And who could foresee what it would take to train a pilot?

Well, this continuing rapid progress in computing and communications hardware means that the big and expensive systems of today will be common in 10 years. Thus, we can learn about the commodity systems of a decade hence by developing and using cutting edge systems today. That is one basis for the emphasis on high performance in the High Performance Computing and Communications Initiative. The fact that the pace of technology is so fast means that only by looking at the very highest performance of today's systems can you anticipate what is going to happen in the very near future.

These very high performance systems give us a window into the future. They are a sort of time machine whereby we can see more clearly and test the medical, military, commercial, entertainment, and economic value of machines that will become commonplace with passing time.

Using today's very fastest systems, we can explore uses not yet practical. Such exploration lets us anticipate and shape our future. It lets us discover the applications, prepare the software programs, and explore uses for systems that will become commonplace soon. I remind you that information systems are only as good as the applications and software programs that run in them, and remember that it takes time, often many years, to refine such applications

and software. The time machine factor afforded by the Federal investment in high performance computing and communications has been very important and will continue to be important to rapid progress of this sector of our economy.

A major success of the HPCCI has been its focus on the task of using parallel computers. It is in fact easier to build lots of transistors than it is to build faster ones, and so we make very powerful computers by running lots of little computers in parallel. It is like trying to harness a team of 1,000 horses. Until such parallel machines existed, no one took seriously the task of writing programs for them, and without programs to run on them, no commercial firm could sensibly offer such high performance machines.

HPCCI broke through this chicken and egg problem by making very fast parallel machines available to advanced research groups much sooner than might otherwise have happened. As a result, programmers demonstrated how to harness their power to important problems in weather prediction, mechanical design, health care, chemistry, and a host of other problems too big to handle otherwise. HPCCI gave U.S. industry the confidence that such machines are useful, and today nearly every computer manufacturer offers some sort of parallel computer. Tomorrow individual silicon chips will have parallel computers. HPCCI built the time machine that revealed this future.

Now, the National Research Council likes me to separate the deliberations of the committee which are in this report from my own personal remarks, and I would like to make a few personal remarks on the subject of why it is that U.S. industry does not look very far forward. There has been, I think, widespread recognition of the fact that U.S. industry is very good at doing development and near-term activities, producing amazing products at short times, but has a rather appalling track record in terms of looking into the farther future. I think there are three reasons for this.

The first one of them involves the risk of long-term research activities. Any one long-term project runs a great risk of failure, and even the biggest companies today cannot afford to fund enough long-term projects to ensure ultimate success. So, managements, I think quite wisely, seek to invest in shorter-term projects.

Nationally, of course, we are big enough to invest in a long-term research program and as we show in our HPCCI report, many of the results of that program that turn out to be the most useful were unanticipated when the project started.

Now, my personal opinion of a second reason involves our mobile society. I work in Silicon Valley where it is endemic that folk change companies fairly regularly. It is the way you get ahead in Silicon Valley, is you go work for somebody else who has a more interesting project or is willing to pay you a little more or has a better stock option situation. That mobile society I think makes it difficult for U.S. industry to invest in long-term research because managements, I think quite correctly, reason how can I invest in teaching folk what the future is going to be like if they are going to go work for somebody else when the payoff comes?

From a national perspective, we have quite a different situation. We do not care which company people work for, and it is fairly rare for U.S. citizens to choose to go and live somewhere else. So, that

knowledge that we build by doing the long-term investments on the national level stays here in the United States where we can in fact make use of it.

My personal opinion of the third reason for the short-term horizon of U.S. industry has to do with U.S. shareholders and financial managers who seem more interested in today's profits than in long-term potential. They cannot evaluate and offer little reward for investment in long-term research. Moreover, they themselves have no incentive to stay with any one company. Rather than forcing the companies that they own to stay abreast of technology, they sell their holdings and move to other companies that have managed somehow to do a better job of anticipating the future.

On the other hand, it is also my personal opinion that the United States is here to stay, let us all hope, for a long time. If the United States does poorly, we cannot sell our holdings in the United States and buy something else. As a nation we can and should make wise long-term investments on which to build our future industry. The rapid pace of information technology fueled by ever improving integrated circuits is a worldwide phenomenon. Others can see the potential of new technology as well as we do. The future strength of our domestic computing and communications industry depends on the present Federal investment in research.

I have now finished my personal remarks. Let me return to our study.

In our study our committee offers several recommendations. By far, the two most important, which we listed as first and second, say in effect, keep the intellectual seed corn alive.

Recommendation one is continue to support research in information technology. Ensure that the major funding agencies, especially the National Science Foundation and the Advanced Research Projects Agency, have strong programs for computing and communications research.

Recommendation two, continue the High Performance Computing and Communications Initiative, maintaining today's increased emphasis on the research challenges posed by the Nation's evolving information infrastructure.

We have a number of other recommendations which flesh out what such a research program might be emphasizing better software techniques for using parallel machines and seeking a firm foundation of research for the communications and information infrastructure that our Nation's economy and defense will need in the coming decade.

We have praise for the interagency cooperation fostered by HPCCI. We identified the need for a full-time coordinator, a post now ably filled by John Toole, but we cautioned that his role should be that of coordination rather than direction, so as to retain the diversity of purpose and the diversity of ideas that characterize a healthy research program. No one individual or small group can anticipate accurately the amazing future that we face. So, it is important that the interagency, broad-scale aspects of this activity be preserved. Only by using the efforts of many independent thinkers can we hope to find the ideas that will fill tomorrow's national needs.

On behalf of our NRC Committee, I thank you for this opportunity to report our findings. Having looked carefully at the HPCCI, we remain strong believers in its importance to our Nation's future.

[The prepared statement of Mr. Sutherland follows:]

**EVOLVING THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS INITIATIVE
TO SUPPORT THE NATION'S INFORMATION INFRASTRUCTURE**

Statement of

**Ivan E. Sutherland, Ph.D.
Co-Chair, Committee on High Performance Computing
and Communications: Status of a Major Initiative
National Research Council
and
Vice President and Fellow, Sun Microsystems**

**before the
Subcommittee on Basic Research
Committee on Science
U.S. House of Representatives**

**hearing on the
High Performance Computing and Communications Program**

October 31, 1995

Good morning, Mr. Chairman and members of the Committee. My name is Ivan Sutherland. I am Vice President and Fellow at Sun Microsystems, a leading manufacturer of computer workstations.

In 1994 Congress asked the National Research Council (NRC) to examine the status of the High Performance Computing and Communications Initiative (HPCCI). I helped respond to that request by serving as co-chair, with Fred Brooks, of the National Research Council's Committee to Study High Performance Computing and Communications: Status of a Major Initiative. We delivered our report early this year. I have copies of our report here for your use, and the Executive Summary is attached to my testimony. The two charts I will use today come from our report.

The National Research Council (NRC) is the operating arm of the National Academy of Sciences (NAS), the National Academy of Engineering (NAE), and the Institute of Medicine (IOM). The National Academy of Sciences was chartered by Congress in 1863 to advise the government on matters of science and technology. I am a member of both the National Academy of Sciences and the National Academy of Engineering.

Let me turn to my first chart. It shows the increase in processing power of information machines over the past 5 years. When an engineer sees such a chart, plotted on a logarithmic scale, he says WOW, because this chart shows amazing progress. What it shows is really just a representation of what we have all experienced. We can all remember the days before FAX; now cellular telephones are ubiquitous. Many of us remember when television was only black and white; now live video in color from abroad appears routinely in our living rooms. We have all observed the explosive growth of the Internet. The power of yesterday's super computer is now in our homes educating and entertaining our children. For a few hundred dollars we can buy a receiver for the global

positioning system that will tell us with amazing precision where we are anywhere on earth.

What is truly amazing about information machines is how quickly they get very much better—about 1000 times better every 15 years, as the chart shows. To understand that, think about the pioneer's covered wagon that went 20 miles a day, early automobiles that went 20 miles an hour, and supersonic aircraft that go 20 miles a minute. Imagine that amount of progress every 15 years, and then realize that built on top of that amazing past progress there remains at least that much more progress yet to come.

Over the past 30 years ever better silicon integrated circuits have driven this revolution in computing and communications. Integrated circuits have gone from single transistors to chips containing tens of millions of transistors in remarkably few years, and as transistors get smaller they go faster. The rapid advance of integrated circuits is the fuel driving the information system revolution. I anticipate improved integrated circuits will continue to fuel this revolution for at least another decade.

The rapid pace of information technology presents great opportunity. The nation that properly anticipates how best to use tomorrow's computing and communications machines gains great advantage in defense, commerce, medicine, education, and quality of life. Ever since World War II we in the United States have enjoyed an advantage in computing and communications. Americans have used the rapid pace of information technology to produce new products, provide new services, form new companies, create new jobs, enhance our defense, and provide important export opportunities. Our ability to use the new devices quickly and well rests in part on our "Yankee ingenuity" and in part on the careful preparation offered by enlightened federal support of long term research.

Our HPCCI report from the NRC contains another chart (figure ES.1) to which I would now like to turn. This chart covers thirty years' progress in some aspects of information technology. On it we show nine of today's major information industry sectors, each of which grew out of federal support of research in computing and communications. In each case we show the interplay between federal support in universities and government laboratories and private industry support in industrial laboratories. The narrow solid lines represent federal support; the diagonally striped lines represent private industry support.

In each case it took 10 years or more for the new ideas to reach market acceptability. We use the bold dotted lines to indicate early commercial product offerings. It took an additional 5 years or so for the combined revenue of these markets to reach a billion dollars annually, which we show as racing stripes. In many cases brand new companies formed around the new ideas from federally supported research to bring the new products to market. As a nation we get our return on federal research investment from the jobs, taxes, and exports produced by these companies as well as from their new products.

I am pleased to work for one such company, Sun Microsystems, a major manufacturer of work stations. Sun is only 13 years old and has revenues of about \$6 billion per year. Sun was started by young people who saw the future of computing and communications in federally supported research programs at Stanford and Berkeley. Sun appears on the chart in the workstation line.

Given the amazingly rapid pace of progress in information systems, how can we anticipate where information technology will lead 5 or 15 years hence? By then, machines with the power of today's super computers will appear in private homes. Perhaps such

machines will understand speech, perhaps they will recognize faces, perhaps they will anticipate human needs far better than we can now imagine. We do not know exactly how, but we do know that better computing and communications will pervade our society. The value of information machines 1000 times more powerful than we have today will be enormous, but they will do things that we do not dream of today. Who in a covered wagon anticipated flying coast to coast in half a day or could foresee what it would take to train a pilot?

Continuing rapid progress in computing and communications hardware means that the big and expensive systems of today will be common in ten years. Thus we can learn about the commodity systems of a decade hence by developing and using cutting edge systems today. That is one basis for the "High Performance" part of HPCCI's name. The very highest performance information systems we can make today give us a window into the future; they are a "time machine" whereby we can see more clearly and test the military, medical, commercial, entertainment and economic value of machines that will become commonplace with passing time.

Using today's very fastest systems we can explore uses not yet practical. Such exploration lets us anticipate and shape our future. It lets us discover the applications, prepare the software programs, and explore the uses for systems that will become commonplace. I remind you that information systems are only as good as the applications and software programs that run in them. Please remember that it takes time, often many years, to refine such applications and software. The "time machine" factor has been and will continue to be essential to rapid progress.

A major success of the HPCCI has been its focus on the task of using parallel computers. It is easier to build lots of transistors than to build faster ones, and so we

make very powerful computers by running lots of little ones in parallel, like a team of 1000 horses. Until such parallel machines existed, no one took seriously the task of writing programs for them, and without programs to run on them no commercial firm would build them.

HPCCI broke through this "chicken and egg" problem by making very fast parallel machines available to advanced research groups much sooner than might otherwise have happened. As a result, programmers demonstrated how to harness their power to important problems in weather prediction, mechanical design, health care, chemistry, and a host of other problems too big to handle otherwise. HPCCI gave U.S. industry the confidence that such machines are useful. Today nearly every computer manufacturer offers some sort of parallel computer. Tomorrow, individual silicon chips will hold parallel computers. HPCCI built the "time machine" that reveals this future.

Left to itself, as is well known, U.S. industry routinely looks forward only a few years. Finding reasons for this was outside the scope of our Committee's work. I would, however, like to offer my personal opinion of three reasons why U.S. industry, quite properly I think, avoids long term research.

My personal opinion of the first reason involves risk. Any one long term project runs great risk of failure. Even the biggest companies today cannot afford to fund enough long term projects to ensure ultimate success. Thus commercial managements, wisely I think, seek to invest in projects that will pay off in a few years rather than a decade. Nationally, however, we are big enough. Many of the fruits of HPCCI have come from unanticipated benefits of broad based exploration possible only in a national program.

Table 1.1 on page 17 of our report lists many examples of such "Unanticipated Results."

My personal opinion of the second reason involves our mobile society. Employee turnover in high technology companies makes it difficult to reap the fruits of long term projects. Why, correctly reasons an executive, should my company support people to work on long term projects when those people may work for my competitor when the projects pay off? From a national perspective, as long as the United States remains a desirable place to live, those same people will surely remain in the United States. From our national perspective we do not care for whom they work so long as they contribute to our domestic economy.

My personal opinion of the third reason involves shareholders. U.S. shareholders and financial managers seem more interested in today's profits than in long term potential. They cannot evaluate and offer little reward for investment in long-term research. They have no incentive to stay with any one company. Rather than forcing their companies to stay abreast of technology, they sell holdings in companies that do not, instead buying stock in newer firms that demonstrate a better grasp of modern ideas.

On the other hand, it is also my personal opinion that our nation is here to stay, let us all hope for a long time. If the United States does poorly we cannot sell our holdings and buy something else. As a nation we can and should make wise investments in long term projects on which to build our future industry. The rapid pace of information technology, fueled by ever improving integrated circuits, is a world-wide phenomenon. Others can see the potential of new technology as well as we do. The future strength of our domestic computing and communications industry depends on present federal investment in research.

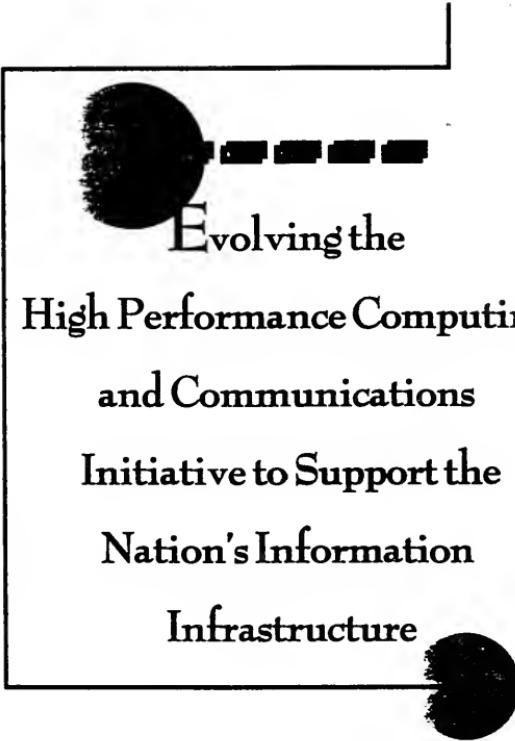
Having finished with my personal opinions about industrial behavior, I return to our NRC report. In it our committee offers several recommendations. By far the two most important say, in effect, "keep the intellectual seed corn alive":

Recommendation 1) Continue to support research in information technology. Ensure that the major funding agencies, especially the National Science Foundation and the Advanced Research Projects Agency, have strong programs for computing and communications research.

Recommendation 2) Continue the High Performance Computing and Communications Initiative, maintaining today's increased emphasis on the research challenges posed by the nation's evolving information infrastructure.

Our other recommendations flesh out what such a research program should be, emphasizing better software techniques for using parallel machines, and seeking a firm foundation of research for the communications and information infrastructure that our nation's economy and defense will need in the coming decades. We have praise for the inter-agency cooperation fostered by the HPCCI. We identified the need for a full time coordinator, a post now ably filled by John Toole, who is here today, but we cautioned that his role should be that of coordination rather than direction, so as to retain the diversity of purpose and diversity of ideas that characterize a healthy research program. No one individual or small group can anticipate accurately the amazing future we face. Only by using the efforts of many independent thinkers can we hope to find the ideas that will fill tomorrow's national needs.

On behalf of our NRC Committee I thank you for this opportunity to report our findings. Having looked carefully at the HPCCI, we remain strong believers in its importance to our nation's future.



Evolving the
High Performance Computing
and Communications
Initiative to Support the
Nation's Information
Infrastructure

Committee to Study High Performance Computing and Communications:
Status of a Major Initiative

Computer Science and Telecommunications Board

Commission on Physical Sciences, Mathematics, and Applications

National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C. 1995

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This report has been reviewed by a group other than the authors according to procedures approved by a Report Review Committee consisting of members of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

The National Academy of Sciences is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce Alberts is president of the National Academy of Sciences.

The National Academy of Engineering was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Robert M. White is president of the National Academy of Engineering.

The Institute of Medicine was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The National Research Council was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce Alberts and Dr. Robert M. White are chairman and vice chairman, respectively, of the National Research Council.

Support for this project was provided by the U.S. Department of Defense through the Advanced Research Projects Agency under Grant No. MDA972-94-1-0008. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Department of Defense or the Advanced Research Projects Agency.

Library of Congress Catalog Card Number 95-67707
International Standard Book Number 0-309-05277-7

Copyright 1995 by the National Academy of Sciences. All rights reserved.

Available from:

National Academy Press
2101 Constitution Avenue, NW
Washington, DC 20418

B-540

Printed in the United States of America

**COMMITTEE TO STUDY
HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS:
STATUS OF A MAJOR INITIATIVE**

FREDERICK P. BROOKS, JR., University of North Carolina at Chapel Hill, *Co-chair*
IVAN E. SUTHERLAND, Sun Microsystems Laboratories, *Co-chair*
ERICH BLOCH, Council on Competitiveness
DEBORAH ESTRIN, University of Southern California/Information Sciences Institute
JOHN HENNESSY, Stanford University
BUTLER W. LAMPSON, Digital Equipment Corporation
EDWARD D. LAZOWSKA, University of Washington
WILLIAM A. LESTER, JR., University of California at Berkeley
JANE PRESTON, Telemedical Interactive Consultative Services Inc.
W. DAVID SINCOSKIE, Bell Communications Research Inc.
LARRY SMARR, National Center for Supercomputing Applications/University of Illinois at Urbana-Champaign
JOSEPH F. TRAUB, Columbia University

Staff

MARJORY S. BLUMENTHAL, Director
JAMES E. MALLORY, Staff Officer
LESLIE M. WADE, Project Assistant

COMPUTER SCIENCE AND TELECOMMUNICATIONS BOARD

WILLIAM WULF, University of Virginia, *Chair*
FRANCES ALLEN, IBM T.J. Watson Research Center
JEFF DOZIER, University of California at Santa Barbara
DAVID J. FARBER, University of Pennsylvania
HENRY FUCHS, University of North Carolina
CHARLES M. GESCHKE, Adobe Systems Inc.
JAMES GRAY, San Francisco, California
BARBARA J. GROSZ, Harvard University
DEBORAH A. JOSEPH, University of Wisconsin
RICHARD M. KARP, University of California at Berkeley
BUTLER W. LAMPSON, Digital Equipment Corporation
BARBARA H. LISKOV, Massachusetts Institute of Technology
JOHN MAJOR, Motorola Inc.
ROBERT L. MARTIN, AT&T Network Systems
DAVID G. MESSERSCHMITT, University of California at Berkeley
WILLIAM H. PRESS, Harvard University
CHARLES L. SEITZ, Myricom Inc.
EDWARD SHORTLIFFE, Stanford University School of Medicine
CASMIRO S. SKRZYPczAK, NYNEX Corporation
LESLIE L. VADASZ, Intel Corporation

MARJORY S. BLUMENTHAL, Director
LOUISE A. ARNHEIM, Senior Staff Officer
HERBERT S. LIN, Senior Staff Officer
JAMES E. MALLORY, Staff Officer
RENEE A. HAWKINS, Staff Associate
JOHN M. GODFREY, Research Associate
GLORIA P. BEMAH, Administrative Assistant
LESLIE M. WADE, Project Assistant

Preface

In early 1994, acting through the Defense Authorization Act for FY 1994 (Public Law 103-160), Congress asked the National Research Council (NRC) to examine the status of the High Performance Computing and Communications Initiative (HPCCI). Broad-based interest in and support for the HPCCI exist. Given its scope and size, concerns had been raised about its goals, management, and progress. Congress asked that at a minimum the study address:

- The basic underlying rationale(s) for the program, including the appropriate balance between federal efforts and private-sector efforts;
- The appropriateness of its goals and directions;
- The balance between various elements of the program;
- The effectiveness of the mechanisms for obtaining the views of industry and users for the planning and implementation of the program;
- The likelihood that the various goals of the program will be achieved;
- The management and coordination of the program; and
- The relationship of the program to other federal support of high-performance computing and communications, including acquisition of high-performance computers by federal departments and agencies in support of the mission needs of these departments and agencies.

For this study the NRC's Computer Science and Telecommunications Board (CSTB) convened a committee of 12 members, expert on pioneering applications of computers and communications and the major components of the HPCCI: High-Performance Computing Systems, Advanced Software Technology and Algorithms, the National Research and Education Network, Basic Research and Human Resources, and Information Infrastructure Technology and Applications. Congress asked the committee to accelerate the normal NRC study process in order to provide an interim report by July 1, 1994, and a final report by February 1, 1995. The committee was able to meet this rapid turnaround by drawing on the knowledge and experience of its members as expressed in committee deliberations and by obtaining input from numerous outside experts.

The full committee met six times between March 10, 1994, and December 20, 1994, to hear more than 25 high-performance computing and communications users, builders, and scientists; to discuss the HPCCI in detail; and to produce this report. Additionally, smaller groups of committee members made site visits to discuss first-hand the use of high-performance technologies. These visits involved another 6 individuals from the Ford Motor Company and approximately 50 high-performance computing and communications users who had gathered at a workshop to discuss the use of high-performance systems in environmental research and simulation. In addition to examining the current status of the program, the committee considered the evolution of the HPCCI

and its goals and alternate government investment strategies related to technological development. The committee took into account varying perspectives on the initiative's goals and available assessments of progress toward achieving them.

The committee's interim report provided technical background and perspective on the overall development of high-performance computing and communications systems, as well as on the HPCCI, formally started in 1991.¹ The interim report made two recommendations: (1) strengthen the National Coordination Office to help it meet increasing future demands for program coordination and information functions; and (2) immediately appoint the congressionally mandated HPCCI Advisory Committee to provide broad-based, active input to the initiative. As this final report goes to press, both recommendations have yet to be acted on and thus require executive and legislative attention.

This final report, *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*, purposely adopted a broad perspective so as to examine the HPCCI within the context of the evolving information infrastructure and national economic competitiveness generally. Committee deliberations consistently pointed out the important contributions that computing and communications research have made to the nation's economy, scientific research, national defense, and social fabric. That research has nourished U.S. leadership in information technology goods, services, and applications. Traditionally, the most powerful computers and the fastest networks made many of those contributions. Recently, however, the widespread availability of significant computing and communications capabilities on the nation's desktops and factory floors has also produced many benefits. The broadening and interconnection of more and more computer-based systems call attention to research needs associated with system scale as well as performance.

The Committee to Study High Performance Computing and Communications: Status of a Major Initiative is grateful for the help, encouragement, and hard work of the NRC staff working with us: Marjory Blumenthal, Jim Mallory, Susan Maurizi, and Leslie Wade. Our meetings went smoothly because of their careful preparation. This report came together because of their attention and diligence. They patiently assembled sometimes conflicting text from diverse authors and helped reconcile it with the critique of our reviewers. They searched out facts of importance to our deliberations. Their excellent staff preparation helped us focus on the substance of our task.

Many others also made valuable contributions to the committee. In addition to individuals listed in Appendix F who briefed the committee, the committee appreciates inputs from Sally Howe and Don Austin (National Coordination Office); Bob Borchers, Paul Young, Dick Kaplan, and Robert Voight (National Science Foundation); Eric Cooper (FORE Systems); Stephen Squires (Advanced Research Projects Agency); Sandy MacDonald (National Oceanic and Atmospheric Administration); Robert Bonometti (Office of Science and Technology Policy); and Al Rosenheck (former congressional staffer). It is also grateful to the anonymous reviewers who helped to sharpen and focus the report with their insightful comments. Responsibility for the report, of course, remains with the committee.

¹Computer Science and Telecommunications Board (CSTB), National Research Council. 1994. *Interim Report on the Status of the High Performance Computing and Communications Initiative*. Computer Science and Telecommunications Board, Washington, D.C.

Finally, we want to acknowledge the contributions to our present task of research projects 30 years past. Of course this text was all word processed. Of course the charts were drafted on computers. Of course we used Internet communication nationwide to plan our meetings, share our thoughts, reconcile our differences, and assemble our report. We plugged our portable computers into a local network at each of our meetings, sending drafts to local laser printers. In short, we have partaken fully of the fruits of the HPCCI's precursors. We thank the visionaries of the past for our tools.

Frederick P. Brooks

Ivan E. Sutherland

Co-chairs

Committee to Study High Performance Computing and Communications:

Status of a Major Initiative

Contents

EXECUTIVE SUMMARY	1
1 U.S. LEADERSHIP IN INFORMATION TECHNOLOGY	13
Information Technology Is Central to Our Society, 13	
Information Technology Advances Rapidly, 14	
Retaining Leadership in Information Technology Is Vital to the Nation, 15	
The Federal Investment in Computing Research Has Paid Rich Dividends, 16	
Continued Federal Investment Is Necessary to Sustain Our Lead, 23	
Today the HPCCI Is the Umbrella for Most Government-sponsored Computing and Communications Research, 25	
Notes, 27	
2 THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS INITIATIVE	28
HPCCI: Goals and Emphases, 28	
Basic Objectives, 28	
Expanded Objectives, 30	
HPCCI Accomplishments, 31	
The Issue of Measurement, 31	
Better Computing and Computational Infrastructure, 32	
Increasing Researcher-Developer-User Synergy, 32	
Impact on Mission Agencies, 36	
Five Gigabit Testbed Projects: Collaboration and Impact, 36	
Evolution of HPCCI Goals and Objectives, 37	
Improving the Information Infrastructure, 37	
Evolving Research Directions and Relevance for the Information Infrastructure, 39	
Overall Computing and Communications R&D Planning, 41	
Toward a Better Balance, 42	
Moving Forward—Basic Issues, 42	
Balance of Private and Public Investment, 42	
Coordination Versus Management, 43	
Coordinating Structure, 45	
Budget, 49	
Notes, 49	

3 RECOMMENDATIONS	53
General Recommendations, 53	
Recommendations on High-Performance Computing, 54	
Recommendations on Networking and Information Infrastructure, 57	
Recommendations on the Supercomputer Centers and Grand Challenge Program, 59	
Recommendations on Coordination and Program Management in the HPCCI, 61	
Comments Relating This Report's Recommendations for High-Performance Computing and Communications Research to Administration Priorities, 64	
Notes, 66	
BIBLIOGRAPHY	67
APPENDIXES	
A The High Performance Computing and Communications Initiative: Background, 77	
B High-Performance Communications Technology and Infrastructure, 92	
C Review of the High Performance Computing and Communications Initiative Budget, 97	
D Current High Performance Computing and Communications Initiative Grand Challenge Activities, 105	
E Accomplishments of National Science Foundation Supercomputer Centers, 108	
F Individuals Providing Briefings to the Committee, 118	

Industry, if left to itself, will naturally find its way to the most useful and profitable employment. Whence it is inferred that manufacturers, without the aid of government, will grow up as soon and as fast as the natural state of things and the interest of the community may require.

Against the solidity of this hypothesis . . . very cogent reasons may be offered . . . [including] the strong influence of habit; the spirit of imitation; the fear of want of success in untried enterprises; the intrinsic difficulties incident to first essays towards [competition with established foreign players]; the bounties, premiums, and other artificial encouragements with which foreign nations second the exertions of their own citizens. . . . To produce the desirable changes as early as may be expedient may therefore require the incitement and patronage of government.

—Alexander Hamilton, 1791, *Report on Manufactures*

Executive Summary

Information technology drives many of today's innovations and offers still greater potential for further innovation in the next decade. It is also the basis for a domestic industry of about \$500 billion,¹ an industry that is critical to our nation's international competitiveness. Our domestic information technology industry is thriving now, based to a large extent on an extraordinary 50-year track record of public research funded by the federal government, creating the ideas and people that have let industry flourish. This record shows that for a dozen major innovations, 10 to 15 years have passed between research and commercial application (see Figure ES.1). Despite many efforts, commercialization has seldom been achieved more quickly.

Publicly funded research in information technology will continue to create important new technologies and industries, some of them unimagined today, and the process will continue to take 10 to 15 years. Without such research there will still be innovation, but the quantity and range of new ideas for U.S. industry to draw from will be greatly diminished. Public research, which creates new opportunities for private industry to use, should not be confused with industrial policy, which chooses firms or industries to support. Industry, with its focus mostly on the near term, cannot take the place of government in supporting the research that will lead to the next decade's advances.

The High Performance Computing and Communications Initiative (HPCCI) is the main vehicle for public research in information technology today and the subject of this report. By the early 1980s, several federal agencies had developed independent programs to advance many of the objectives of what was to become the HPCCI. The program received added impetus and more formal status when Congress passed the High Performance Computing Act of 1991 (Public Law 102-194) authorizing a 5-year program in high-performance computing and communications. The initiative began with a focus on high-speed parallel computing and networking and is now evolving to meet the needs of the nation for widespread use on a large scale as well as for high speed in computation and communications. To advance the nation's information infrastructure there is much that needs to be discovered or invented, because a useful "information highway" is much more than wires to every house.

As a prelude to examining the current status of the HPCCI, this report first describes the rationale for the initiative as an engine of U.S. leadership in information technology and outlines the contributions of ongoing publicly funded research to past and current progress in developing computing and communications technologies (Chapter 1). It then describes and evaluates the HPCCI's goals, accomplishments, management, and planning (Chapter 2). Finally, it makes recommendations aimed at ensuring continuing U.S. leadership in information technology through wise evolution and use of the HPCCI as an important lever (Chapter 3). Appendixes A through F of the report provide additional details on and documentation for points made in the main text.

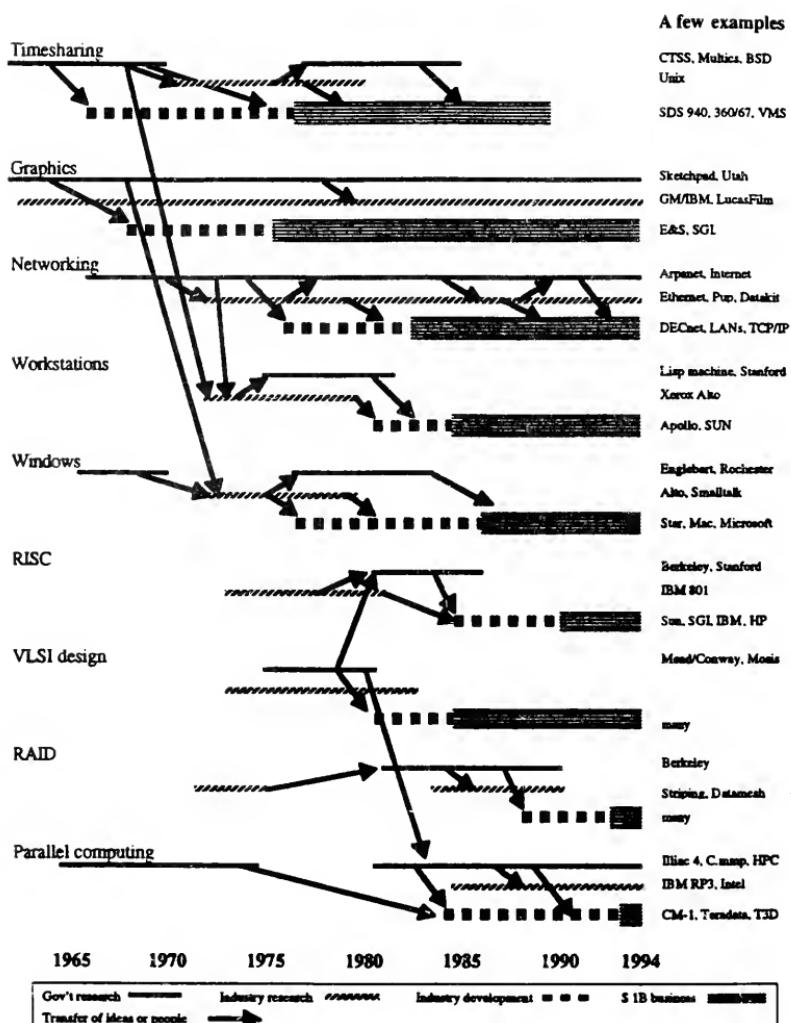


FIGURE ES.1 Government-sponsored computing research and development stimulates creation of innovative ideas and industries. Dates apply to horizontal bars, but not to arrows showing transfer of ideas and people.

INFORMATION TECHNOLOGY—FUNDAMENTAL FOR SOCIETY AND THE ECONOMY NOW AND TOMORROW

Computers, the devices that process information, affect our lives both directly and indirectly. Today, more than 70 million microcomputers are installed in the United States, and between one-fifth and one-third of U.S. households have one.² Entertainment, education, communications, medicine, government, and finance are using computers in more ways to enhance our lives directly through the provision of such services as distributed learning and remote banking. Computers are also used to make essential products and activities cheaper and better: airplanes, molded plastics, automobiles, medical imaging, and oil exploration are only a few of many examples. A broader benefit is the \$500 billion industry's creation of jobs, taxes, profits, and exports.

Clearly, the uses and applications of information technology will continue to grow. In fact, the information revolution has only just begun. Computers will become increasingly valuable to industries and to citizens as their power is tapped to recognize and simulate speech, generate realistic images, provide accurate models of the physical world, build huge automated libraries, control robots, and help with a myriad of other tasks. To do these things well will require both computing and communications systems many times more powerful than we have today. Ongoing advances in knowledge will constitute the foundation for building the systems and developing the applications that will continue to advance our quality of life and ensure strong U.S. leadership in information technology. Strong leadership in information technology in turn supports other sectors including industry, health, education, and defense by serving their needs for equipment, software, and know-how.

The Basis for Continuing Strength— A Successful Government-Industry Partnership

Federal investment in information technology research has played a key role in the U.S. capability to maintain its international lead in information technology. Starting in World War II publicly funded research has helped to stock the nation's storehouse of trained people and innovative ideas. But our lead is fragile. Leadership can shift in a few product generations, and because a generation in the computing and communications industry is at most 2 years, our lead could disappear in less than a decade.

Since the early 1960s the U.S. government has invested broadly in computing research, creating new ideas and trained people. The result has been the development of important new technologies for time-sharing, networking, computer graphics, human-machine interfaces, and parallel computing, as well as major contributions to the design of very large scale integrated circuits, fast computers and disk systems, and workstations (see Figure ES.1; see also Chapter 1, Box 1.2 for details). Each of these is now a multibillion-dollar business. From these successes we can learn some important lessons:

- *Research has kept paying off over a long period.*
- *The payoff from research takes time.* As Figure ES.1 shows, at least 10 years, more often 15, elapse between initial research on a major new idea and commercial success. This is still true in spite of today's shorter product cycles.
- *Unexpected results are often the most important.* Electronic mail and the "windows" interface are only two examples; Box 1.2 in Chapter 1 outlines more.

- *Research stimulates communication and interaction.* Ideas flow back and forth between research programs and development efforts and between academia and industry.
- *Research trains people,* who start companies or form a pool of trained personnel that existing companies can draw on to enter new markets quickly.
- *Doing research involves taking risks.* Not all public research programs have succeeded or led to clear outcomes even after many years. But the record of accomplishments suggests that government investment in computing and communications research has been very productive.

Government Support of Research Is Crucial

The information technology industry improves its products faster than most others: for the last 40 years a dollar has bought hardware with twice as much computation, storage, and communication every 18 to 24 months, offering a 100-fold gain every decade (Patterson and Hennessy, 1994, p. 21). This rate will continue at least for the next decade (see Chapter 1, Figure 1.1). Better hardware in turn makes it feasible to create software for new applications: electronic and mechanical design, climate mapping, digital libraries, desktop publishing, video editing, and telemedicine are just a few examples. Such applications are often brought to market by new companies such as Microsoft and Sun Microsystems, both of which produce revenues of more than \$4 billion per year (Computer Select, 1994) and neither of which existed 15 years ago.

The information technology industry is characterized by great importance to the economy and society, rapid and continuing change, a 10- to 15-year cycle from major idea to commercial success, and successive waves of new companies. In this environment a broad program of publicly funded research is essential for two reasons:

- First, industrial efforts cannot replace government investment in basic computing and communications research. Few companies will invest for a payoff that is 10 years away, and even a company that does make a discovery may postpone using it. The vitality of the information technology industry depends heavily on new companies, but new companies cannot easily afford to do research; furthermore, industry in general is doing less research now than in the recent past (Geppert, 1994; Corcoran, 1994). But because today's sales are based on yesterday's research, investment in innovation must go forward so that the nation's information industry can continue to thrive.
- Second, it is hard to predict which new ideas and approaches will succeed. The exact course of exploratory research cannot be planned in advance, and its progress cannot be measured precisely in the short term. The purpose of publicly funded research is to advance knowledge and create new opportunities that industry can exploit in the medium and long term, not to determine how the market should develop.

THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS INITIATIVE

Goals and Emphases

The HPCCI is the current manifestation of the continuing government research program in information technology, an investment that has been ongoing for more than 50 years. Although it emphasizes research in high-performance computing and communications, the HPCCI now has in its budget nearly all of the federal funding for computing research of any kind. The wisdom of this arrangement is doubtful.

The HPCCI was initiated to serve several broad goals (NCO, 1993):

- Extend U.S. leadership in high-performance computing and networking;
- Disseminate new technologies to serve the economy, national security, education, health care, and the environment; and
- Spur gains in U.S. productivity and industrial competitiveness.

The original plans to achieve these goals called for creating dramatically faster computers and networks, stretching their limits with Grand Challenge problems in scientific computing, setting up supercomputer centers with the machines and experts needed to attack these challenges, and training people to build and exploit the new technology. More recently the focus has been shifting toward broader uses of computing and communications.

High Performance

"High performance"—which involves bringing more powerful computing and communications technology to bear on a problem—has enabled advances on several fronts. High-performance systems, for example, deliver answers sooner for complex problems that need large amounts of computing. Timely and accurate forecasting of weather, mapping of oil reservoirs, and imaging of tumors are among the benefits encompassed by the goals listed above. But "high performance," which is broader than supercomputing, is a moving target because of the steady and rapid gains in the performance/cost ratio. Yesterday's supercomputer is today's personal computer; today's leading-edge communications technology will be among tomorrow's mainstream capabilities.

Information technology evolves as new and valuable applications are found for hardware that gets steadily more powerful and cheaper. To benefit, users need affordable hardware, but they also need the software that implements the new applications. Yet learning how to build software takes many years of experimentation. If this process starts only when the hardware has already become cheap, the benefits to users will be delayed by years. Research needs to treat today's expensive equipment as a time machine, learning how it will be used when it is cheap and widely available, as it surely will be tomorrow. Knowing how to use computers for new tasks sooner can help many industries to become more competitive.

To date, the HPCCI's focus has been mainly on speed, but speed is not the only measure of high performance. Both speed, measured today in billions of operations per second or billions of bits per second, and scale, measured by the number of millions of users served, are important research issues. However, for the nation's information infrastructure, scale now seems more difficult to achieve. Information technology can be thought of as a tent, with the height of the center pole as speed and the breadth of the base as scale. Widening the tent to allow more work on scale without decreasing the work on speed requires more cloth; with the same resources, widening

the tent would sacrifice research on speed for research on scale. This report recommends ways to reallocate funds within the HPCCI so as to accommodate greater emphasis on scale.

Accomplishments to Date

The HPCCI has focused mainly on parallel computing, networking, and development of human resources. Building on progress in research begun before the HPCCI, work and accomplishments to date reveal two key trends: better computing and computational infrastructure and increasing researcher-developer-user synergy.

Despite the difficulty of measuring impact at this early stage, it is the committee's judgment that the HPCCI has been generally successful so far. That assessment is necessarily qualitative and experiential now. Because the HPCCI is only 3 years old, results that can be measured in dollars should not be expected before the next decade.

The HPCCI has contributed substantially to the development, deployment, and understanding of computing and communications facilities and capabilities as infrastructure. It has helped transform understanding of how to share resources and information, generating proofs of concept and understanding that are of value not only to the scientific research community but also to the economy and society at large.

In parallel computing the fundamental challenge is not building the machines, but learning how to program them. Pioneering users and their software developers must be motivated by machines that are good enough to reward success with significant speedups.³ For this reason, a great deal of money and effort have had to be spent to obtain parallel machines with the potential to run much faster than existing supercomputers. From the base built by the HPCCI, much has been learned about parallel computing.

The HPCCI has fostered productive interactions among the researchers and developers involved in creating high-performance computing and communications technology and researchers who use the technology. Building on the varying perspectives of the three groups, complex problems are being solved in unique ways. In particular, the HPCCI has funded cross-disciplinary teams associated with the Grand Challenge projects to solve complex computational problems and produce essential new software for the new parallel systems.

More specifically, the HPCCI has:

- Increased the nation's stock of expertise by educating new students and attracting new researchers;
- Made parallel computing widely accepted as the practical route to achieving high-performance computing;
- Demonstrated the feasibility of and initiated deployment of parallel databases;
- Driven progress on Grand Challenge problems in disciplines such as cosmology, molecular biology, chemistry, and materials science. Parallel computation has enhanced the ability to attack problems of great complexity in science and engineering;
- Developed new modes of analyzing and visualizing complex data sets in the earth sciences, medicine, molecular biology, and engineering, including creating virtual reality technologies. Many supercomputer graphic techniques of the 1980s are now available on desktop graphics workstations;

- Through the gigabit network testbeds associated with the National Research and Education Network component, demonstrated the intimate link between computing and communications systems;
- Built advanced networks that are the backbone of the Internet and the prototypes for its further evolution into the basis for a broader information infrastructure;
- Deployed a high-speed backbone that has kept up with the yearly doubling of the size of the Internet, and organized the impending transition of this backbone away from government funding; and
- Created the Mosaic browser for the World Wide Web, the first new major application in many years that promises to greatly increase access to the resources available on the Internet. This was an entirely unexpected result.

Evolution

A large-scale, integrated information infrastructure designed to serve the entire nation is becoming a high priority for government and industry as well as a source of challenges for research. Complex systems with millions of users pose many problems: performance, management, security, interoperability, compatible evolution of components, mobility, and reliability are only a few. Today's technology can solve these problems for systems with a few thousand users at most; to do so for millions or hundreds of millions of users is far beyond the current state of the art.⁴ Providing users with high-bandwidth connections is itself a major problem, but it is only the beginning. There is a wide gap between enabling a connection and providing a rich array of useful and dependable services.

Because the HPCCI has become the rubric under which virtually all of the nation's research in information technology is conducted, it is not surprising that its focus has been changing in response to past successes, new opportunities, and evolving societal needs. The recently added Information Infrastructure Technology and Applications (IITA) program, broadly construed, addresses many of the problems just mentioned; it is already the largest component of the HPCCI,⁵ and its continued evolution should be encouraged.

But with the policy focus—in the government, the press, and in most of the agencies—centered on information infrastructure,⁶ high-performance computing seems to have been downplayed. The committee emphasizes the importance of retaining the HPCCI's momentum at just the time when its potential to support improvement in the nation's information infrastructure is most needed.

Organization

Several federal agencies participate in the HPCCI, most notably the National Science Foundation (NSF), the Advanced Research Projects Agency (ARPA), the National Aeronautics and Space Administration (NASA), and the Department of Energy (DOE) (see Appendix A for a full list). Because of its successes the HPCCI has become a model for multiagency collaboration and for the "virtual agency" concept advanced through the *National Performance Review* (Gore, 1993). Each participating agency retains responsibility for its own part of the program, but the agencies work together in joint funding of projects, such as the federal portions of the Internet; joint reviews of grants and contracts, such as the NSF-ARPA-NASA digital library initiative; joint testbeds; and

consortia, such as the consortium for Advanced Modeling of Regional Air Quality that joins six federal agencies with several state and local governments.

The HPCCI supports a diverse set of contractors at universities, companies, and national laboratories throughout the country. It provides project funding in varying amounts through contracts, grants, and cooperative agreements awarded according to diverse methods. This diversity is healthy because it allows many views to compete, resulting in a broad research program that ensures a continuing flow of advances in information technology.

Some have argued for a more centrally managed program, with thorough planning, precise milestones, and presumably no wasted effort. Tighter management would cost more in bureaucracy and turf wars, but the essential question is whether it would produce better or worse results for the money spent. The committee believes that because of the long time scale of research, diversity is essential for success. No one person or organization is either smart or lucky enough to plan the best program, no single approach is best, and success often comes in unanticipated ways. Because it is a national research program and because of the many different but interdependent underlying technologies, the HPCCI is necessarily and properly far more diverse than a focused effort such as the Apollo moon landing program or a commercial product development program.

In contrast to central management, coordination enhances the benefits of diversity by helping to prevent unintended duplication, redundancy, and missed opportunities. The HPCCI's National Coordination Office (NCO) serves this purpose, aiding interagency cooperation and acting as liaison for the initiative to the Congress, other levels of government, universities, industry, and the public. Its efforts are reflected in its impressive FY 1994 and FY 1995 "Blue Books" describing the program's activities and spending.⁷ Strengthening the NCO and appointing an advisory committee, as recommended in the committee's interim report (CSTB, 1994c), would facilitate regular infusions of ideas and advice from industry and academia and enable better communication of the HPCCI's goals and accomplishments to its many constituents. This committee should consist of a group of recognized experts that is balanced between academia and industry and balanced with respect to application areas and the core technologies underlying the HPCCI.

Budget

Because it grew from earlier programs, a significant portion of the HPCCI budget is not new money. The budget grew from a base of \$490 million in preexisting funding in FY 1992 to the \$1.1 billion requested for FY 1995.⁸ Each year agencies have added to the base by moving budgets for existing programs into the HPCCI and by reprogramming existing funds to support the HPCCI. Congress has also added funding each year to start new activities or expand old ones.

The result is that much of the \$1.1 billion requested for FY 1995 is money that was already being spent on computing and communications in FY 1992. The request has three elements: (1) funds for activities that predate the HPCCI and were in the FY 1992 base budget, (2) funds for activities that have since been designated as part of the HPCCI, and (3) new funds for new activities or for growth. Although dissecting the budget in this way would shed light on the program, the committee was unable to do so because each participating agency treats the numbers differently.

It appears that the FY 1995 request breaks down roughly as one-third for applications, one-third to advance the essential underlying computing and communications technologies, one-quarter for computing and communications infrastructure, and small amounts for education and electronics (see Appendix C).

THE FUTURE OF THE HPCCI: RECOMMENDATIONS

The committee believes that strong public support for a broadly based research program in information technology is vital to maintaining U.S. leadership in information technology. Incorporating this view of the importance and success of the government's investment in research, the 13 recommendations that follow address five areas: general research program, high-performance computing, networking and information infrastructure, the supercomputer centers and the Grand Challenge projects, and program coordination and management. Within each area the recommendations are presented in priority order.

General Recommendations

1. Continue to support research in information technology. Ensure that the major funding agencies, especially the National Science Foundation and the Advanced Research Projects Agency, have strong programs for computing and communications research that are independent of any special initiatives.

The government investment in computing research has yielded significant returns. Ongoing investment, *at least as high as the current dollar level*, is critical both to U.S. leadership and to ongoing innovation in information technology. Today the HPCCI supports nearly all of this research, an arrangement that is both misleading and dangerous: misleading because much important computing research addresses areas other than high performance (even though it may legitimately fit under the new IITA component of the HPCCI), and dangerous because reduced funding for the HPCCI could cripple all of computing research. The "war on cancer" did not support all of biomedical research, and neither should the HPCCI or any future initiative on national infrastructure subsume all of computing research.

2. Continue the HPCCI, maintaining today's increased emphasis on the research challenges posed by the nation's evolving information infrastructure. The new Information Infrastructure Technology and Applications program of the HPCCI focuses on information infrastructure topics, which are also addressed in the initiative's other four components. The committee supports this continued evolution, which will lead to tangible returns on existing and future investments in basic hardware, networking, and software technologies.

High-Performance Computing

3. Continue funding a strong experimental research program in software and algorithms for parallel computing machines. Today a crucial obstacle to widespread use of parallel computing is the lack of advanced software and algorithms. Emphasis should be given to research on developing and building usable applications-oriented software systems for parallel computers. Avoid funding the transfer ("porting") of existing commercial applications to new parallel computing machines unless there is a specific research need.

4. Stop direct HPCCI funding for development of commercial hardware by computer vendors and for "industrial stimulus" purchases of hardware. Maintain HPCCI support for precompetitive research in computer architecture; this work should be done in universities or in university-industry collaborations and should be driven by the needs of system and application software. HPCCI funding for stimulus purchase of large-scale machines has been reduced, as has the funding of hardware development by vendors. The committee supports these changes, which should continue except when a mission need demands the development of nonstandard hardware.

Public research is best done in universities. Not only are academic organizations free to think about longer-term issues, but they also stimulate technology transfer through publication and placement of graduates. The national experience supports Vannevar Bush's basic tenet: publicly funded research carried out in universities produces a diversity of excellent ideas, trained people, research results, and technologies that can be commercially exploited (OSRD and Bush, 1945).

5. Treat development of a teraflop computer as a research direction rather than a destination. The goal of developing teraflop capability has served a valuable purpose in stimulating work on large-scale parallelism, but further investment in raw scalability is inappropriate except as a focus for precompetitive, academic research. Industrial development of parallel computers will balance the low cost of individual, mass-produced computing devices against the higher cost of communicating between them in a variety of interesting ways. In the near future a teraflop parallel machine will be built when some agencies' mission requirements correspond to a sufficiently economical commercial offering. Continued progress will surely lead to machines even larger than a teraflop.

Networking and Information Infrastructure

New ideas are needed to meet the new challenges underlying development of the nation's information infrastructure. The HPCCI can contribute most by focusing on the underlying research issues. This shift has already begun, and it should continue.

This evolution of the research agenda, which would support improvement of the nation's information infrastructure, is partly under way: in the *FY 1995 Implementation Plan* (NCO, 1994, p. 15), over one-quarter of the NSF and ARPA HPCCI funding is focused on the IITA component, and activities in other components have also evolved consistent with these concerns. The committee supports this increased emphasis.

6. Increase the HPCCI focus on communications and networking research, especially on the challenges inherent in scale and physical distribution. An integrated information infrastructure that fully serves the nation's needs cannot spring full-grown from what we already know. Much research is needed on difficult problems related to size, evolution, introduction of new systems, reliability, and interoperability. Much more is involved than simply deploying large numbers of boxes and wires. For example, both hardware and software systems must work efficiently to handle scheduling; bandwidth optimization for transmission of a range of data formats, including real-time audio and video data; protocol and format conversion; security; and many other requirements.

7. Develop a research program to address the research challenges underlying our ability to build very large, reliable, high-performance, distributed information systems based on the existing HPCCI foundation. Although a comprehensive vision of the research needed for advancing the nation's information infrastructure has not yet been developed, three key areas for research are scalability, physical distribution, and interoperative applications.

8. Ensure that research programs focusing on the National Challenges contribute to the development of information infrastructure technologies as well as to the development of new applications and paradigms. This dual emphasis contrasts with the narrower focus on scientific results that has driven work on the Grand Challenges.

Supercomputer Centers and Grand Challenge Program

The NSF supercomputer centers have played a major role in establishing parallel computing as a full partner with the prior paradigms of scalar and vector computing by providing access to

state-of-the-art computing facilities. NSF should continue to take a broad view of the centers' mission of providing access to high-performance computing and communications resources, including participating in research needed to improve software for parallel machines and to advance the nation's information infrastructure.

The committee recognizes that advanced computation is an important tool for scientists and engineers and that support for adequate computer access must be a part of the NSF research program in all disciplines. The committee did not consider the appropriate overall funding level for the centers. However, the committee believes that NSF should move to a model similar to that used by NASA and DOE for funding general access to computing. The committee prefers NASA's and DOE's approach to funding supercomputer centers, where HPCCI funds are used only to support the exploration and use of new computing architectures, while non-HPCCI funds are used to support general access.

9. The mission of the National Science Foundation supercomputer centers remains important, but the NSF should continue to evaluate new directions, alternative funding mechanisms, new administrative structures, and the overall program level of the centers. NSF could continue funding of the centers at the current level or alter that level, but it should continue using HPCCI funds to support applications that contribute to the evolution of the underlying computing and communications technologies, while support for general access by application scientists to maturing architectures should come increasingly from non-HPCCI funds.

10. The Grand Challenge program is an innovative approach to creating interdisciplinary and multi-institutional scientific research teams; however, continued use of HPCCI funds is appropriate only when the research contributes significantly to the development of new high-performance computing and communications hardware or software. Grand Challenge projects funded under the HPCCI should be evaluated on the basis of their contributions both to high-performance computing and communications technologies and to the application area. Completion of the Grand Challenge projects will provide valuable insights and demonstrate the capabilities of new high-performance architectures in some important applications. It will also foster better collaboration between computer scientists and computational scientists. The committee notes that a large share of HPCCI funding for the Grand Challenges currently comes from the scientific disciplines involved. However, the overall funding seems to come entirely from HPCCI-labeled funds. For the same reasons outlined in Recommendation 9, the committee sees this commingled support as unhealthy in the long run and urges a transition to greater reliance on scientific disciplinary funding using non-HPCCI funds.

Coordination and Program Management

11. Strengthen the HPCCI National Coordination Office while retaining the cooperative structure of the HPCCI and increasing the opportunity for external input. Immediately appoint the congressionally mandated advisory committee intended to provide broad-based, active input to the HPCCI, or provide an effective alternative. Appoint an individual to be a full-time coordinator, program spokesperson, and advocate for the HPCCI.

In making this recommendation, the committee strongly endorses the role of the current NCO as supporting the mission agencies rather than directing them. The committee believes it vital that the separate agencies retain direction of their HPCCI funds. The value of interagency cooperation outweighs any benefits that might be gained through more centralized management.

Diverse management systems for research should be welcomed, and micromanagement should be avoided. In the past, choosing good program officers and giving them freedom to operate independently have yielded good value, and the committee believes it will continue to do so.

Furthermore, independence will encourage diversity in the research program, thus increasing opportunities for unexpected discoveries, encouraging a broader attack on problems, and ensuring fewer missed opportunities.

12. Place projects in the HPCCI only if they match well to its objectives. Federal research funding agencies should promptly document the extent to which HPCCI funding is supporting important long-term research areas whose future funding should be independent of the future of the HPCCI.

A number of preexisting agency programs have entered the HPCCI, with two effects: the HPCCI's budget appears to grow faster than the real growth of investment in high-performance computing and communications research, and important programs such as basic research in computing within NSF and ARPA may be in jeopardy should the HPCCI end.

13. Base mission agency computer procurements on mission needs only, and encourage making equipment procurement decisions at the lowest practical management level. This recommendation applies equally to government agencies and to government contractors. It has generally been best for an agency to specify the results it wants and to leave the choice of specific equipment to the contractor or local laboratory management.

NOTES

1. See U.S. DOC (1994); the Department of Commerce utilizes data from the U.S. Bureau of the Census series, the *Annual Survey of Manufactures*. It places the value of shipments for the information technology industry at \$421 billion for 1993. This number omits revenue from equipment rentals, fees for after-sale service, and mark-ups in the product distribution channel. It also excludes office equipment in total. It includes computers, storage devices, terminals and peripherals; packaged software; computer program manufacturing, data processing, information services, facilities management, and other services; and telecommunications equipment and services.

See also CBEMA (1994); CBEMA values the worldwide 1993 revenue of the U.S. information technology industry at \$602 billion. In addition to including office equipment, it shows larger revenues for information technology hardware and telecommunications equipment than does the Department of Commerce.

2. Microcomputers (personal computers) are defined as computers with a list price of \$1,000 to \$14,999; see CBEMA (1994), pp. 60-61. Forrester Research Inc. (1994, pp. 2-3) estimates the share of households with PCs at about 20 percent, based on its survey of households and Bureau of Census data. Forrester's model accounts for retirements of older PCs and for households with multiple PCs. This is a lower estimate than the Software Publishing Association's widely cited 30 percent share. By definition, the microcomputer statistics exclude small computers and other general-purpose and specialized devices that also make use of microprocessors and would be counted in a more comprehensive measurement of information technology.

3. Earlier experience with three isolated computers, "Illiac 4" (built at the University of Illinois) and "C.mmp" and "Cm*" (both built at Carnegie Mellon University), bears out this point.

4. Of course, systems specialized for a single application or for homogenous technology, such as telephony, serve millions of users, but what is now envisioned is more complex and heterogenous, involving integration of multiple services and systems.

5. The other four programs of the HPCCI are Advanced Software Technology and Algorithms, Basic Research and Human Resources, High-Performance Computing Systems, and the National Research and Education Network.

6. Notably, references to the computing portion of the HPCCI have been overshadowed recently by the ubiquity of speeches and documents devoted to the notion of a national information infrastructure (NII). The NII has also been featured in the titles of the 1994 and 1995 Blue Books.

7. Each year beginning in 1991 the director of the Office of Science and Technology Policy submits a report on the HPCCI to accompany the president's budget. The FY 1992, FY 1993, and FY 1994 books were produced by the now-defunct Federal Coordinating Council for Science, Engineering, and Technology; the FY 1995 report was produced by the NCO (acting for the Committee on Information and Communications). The report describes prior accomplishments and the future funding and activities for the coming fiscal year. These reports have collectively become known as "Blue Books" after the color of their cover.

8. NCO (1994), p. 15. Note that figures represent the President's requested budget authority for FY 1995. Actual appropriated levels were not available at press time. Because the HPCCI is synthesized as a cross-cutting multiagency initiative, there is no separate and identifiable "HPCCI appropriation."

*Evolving the HPPCI to Support the
Nation's Information Infrastructure, p. 14*

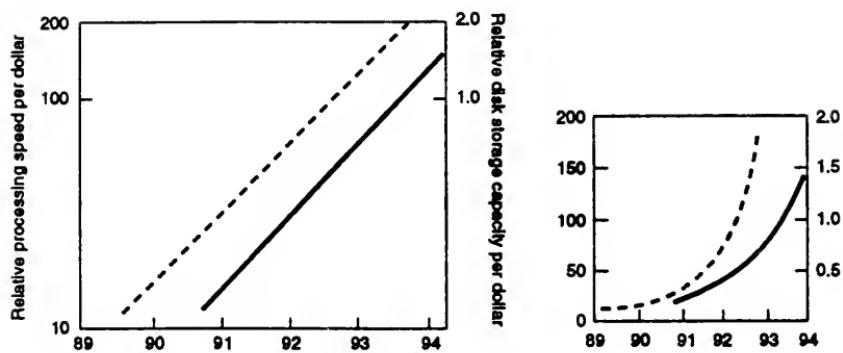
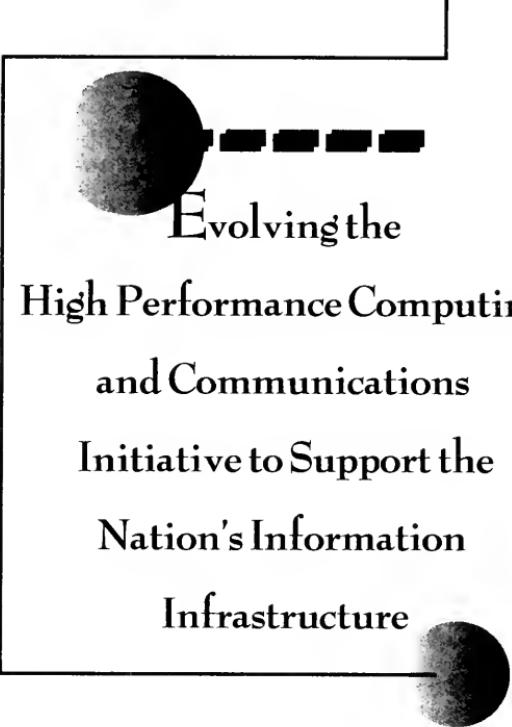


FIGURE 1.1 Increase in performance per dollar of processor speed and disk storage from 1989 to 1994, shown on a semilog scale. (The right-hand graph uses a linear scale to emphasize the compound effect of successive doublings.)

Mr. SCHIFF. Thank you very much, Mr. Sutherland.

I would first like to ask unanimous consent that the report that Mr. Sutherland referred to entitled Evolving the High Performance Computing and Communications Initiative—Support the Nation's Information Infrastructure, published by the National Research Council be made a part of this record. Without objection, it will be made a part of the record.

[The report referred to follows.]



Evolving the
High Performance Computing
and Communications
Initiative to Support the
Nation's Information
Infrastructure

Committee to Study High Performance Computing and Communications:
Status of a Major Initiative

Computer Science and Telecommunications Board

Commission on Physical Sciences, Mathematics, and Applications

National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C. 1995

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

This report has been reviewed by a group other than the authors according to procedures approved by a Report Review Committee consisting of members of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

The National Academy of Sciences is a private, nonprofit, self-perpetuating society of distinguished scholars engaged in scientific and engineering research, dedicated to the furtherance of science and technology and to their use for the general welfare. Upon the authority of the charter granted to it by the Congress in 1863, the Academy has a mandate that requires it to advise the federal government on scientific and technical matters. Dr. Bruce Alberts is president of the National Academy of Sciences.

The National Academy of Engineering was established in 1964, under the charter of the National Academy of Sciences, as a parallel organization of outstanding engineers. It is autonomous in its administration and in the selection of its members, sharing with the National Academy of Sciences the responsibility for advising the federal government. The National Academy of Engineering also sponsors engineering programs aimed at meeting national needs, encourages education and research, and recognizes the superior achievements of engineers. Dr. Robert M. White is president of the National Academy of Engineering.

The Institute of Medicine was established in 1970 by the National Academy of Sciences to secure the services of eminent members of appropriate professions in the examination of policy matters pertaining to the health of the public. The Institute acts under the responsibility given to the National Academy of Sciences by its congressional charter to be an adviser to the federal government and, upon its own initiative, to identify issues of medical care, research, and education. Dr. Kenneth I. Shine is president of the Institute of Medicine.

The National Research Council was organized by the National Academy of Sciences in 1916 to associate the broad community of science and technology with the Academy's purposes of furthering knowledge and advising the federal government. Functioning in accordance with general policies determined by the Academy, the Council has become the principal operating agency of both the National Academy of Sciences and the National Academy of Engineering in providing services to the government, the public, and the scientific and engineering communities. The Council is administered jointly by both Academies and the Institute of Medicine. Dr. Bruce Alberts and Dr. Robert M. White are chairman and vice chairman, respectively, of the National Research Council.

Support for this project was provided by the U.S. Department of Defense through the Advanced Research Projects Agency under Grant No. MDA972-94-1-0008. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Department of Defense or the Advanced Research Projects Agency.

Library of Congress Catalog Card Number 95-67707
International Standard Book Number 0-309-05277-7

Copyright 1995 by the National Academy of Sciences. All rights reserved.

Available from:

National Academy Press
2101 Constitution Avenue, NW
Washington, DC 20418

B-540

Printed in the United States of America

**COMMITTEE TO STUDY
HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS:
STATUS OF A MAJOR INITIATIVE**

FREDERICK P. BROOKS, JR., University of North Carolina at Chapel Hill, *Co-chair*
IVAN E. SUTHERLAND, Sun Microsystems Laboratories, *Co-chair*
ERICH BLOCH, Council on Competitiveness
DEBORAH ESTRIN, University of Southern California/Information Sciences Institute
JOHN HENNESSY, Stanford University
BUTLER W. LAMPSON, Digital Equipment Corporation
EDWARD D. LAZOWSKA, University of Washington
WILLIAM A. LESTER, JR., University of California at Berkeley
JANE PRESTON, Telemedical Interactive Consultative Services Inc.
W. DAVID SINCOSKIE, Bell Communications Research Inc.
LARRY SMARR, National Center for Supercomputing Applications/University of Illinois at Urbana-Champaign
JOSEPH F. TRAUB, Columbia University

Staff

MARJORY S. BLUMENTHAL, Director
JAMES E. MALLORY, Staff Officer
LESLIE M. WADE, Project Assistant

COMPUTER SCIENCE AND TELECOMMUNICATIONS BOARD

WILLIAM WULF, University of Virginia, *Chair*
FRANCES ALLEN, IBM T.J. Watson Research Center
JEFF DOZIER, University of California at Santa Barbara
DAVID J. FARBER, University of Pennsylvania
HENRY FUCHS, University of North Carolina
CHARLES M. GESCHKE, Adobe Systems Inc.
JAMES GRAY, San Francisco, California
BARBARA J. GROSZ, Harvard University
DEBORAH A. JOSEPH, University of Wisconsin
RICHARD M. KARP, University of California at Berkeley
BUTLER W. LAMPSON, Digital Equipment Corporation
BARBARA H. LISKOV, Massachusetts Institute of Technology
JOHN MAJOR, Motorola Inc.
ROBERT L. MARTIN, AT&T Network Systems
DAVID G. MESSERSCHMITT, University of California at Berkeley
WILLIAM H. PRESS, Harvard University
CHARLES L. SEITZ, Myricom Inc.
EDWARD SHORTLIFFE, Stanford University School of Medicine
CASMIR S. SKRZYPCZAK, NYNEX Corporation
LESLIE L. VADASZ, Intel Corporation

MARJORY S. BLUMENTHAL, Director
LOUISE A. ARNHEIM, Senior Staff Officer
HERBERT S. LIN, Senior Staff Officer
JAMES E. MALLORY, Staff Officer
RENEE A. HAWKINS, Staff Associate
JOHN M. GODFREY, Research Associate
GLORIA P. BEMAH, Administrative Assistant
LESLIE M. WADE, Project Assistant

**COMMISSION ON PHYSICAL SCIENCES,
MATHEMATICS, AND APPLICATIONS**

RICHARD N. ZARE, Stanford University, *Chair*
RICHARD S. NICHOLSON, American Association for the
Advancement of Science, *Vice Chair*
STEPHEN L. ADLER, Institute for Advanced Study
SYLVIA T. CEYER, Massachusetts Institute of Technology
SUSAN L. GRAHAM, University of California at Berkeley
ROBERT J. HERMANN, United Technologies Corporation
RHONDA J. HUGHES, Bryn Mawr College
SHIRLEY A. JACKSON, Rutgers University
KENNETH I. KELLERMANN, National Radio Astronomy Observatory
HANS MARK, University of Texas at Austin
THOMAS A. PRINCE, California Institute of Technology
JEROME SACKS, National Institute of Statistical Sciences
L.E. SCRIVEN, University of Minnesota
A. RICHARD SEEBASS III, University of Colorado
LEON T. SILVER, California Institute of Technology
CHARLES P. SLICHTER, University of Illinois at Urbana-Champaign
ALVIN W. TRIVELPIECE, Oak Ridge National Laboratory
SHMUEL WINOGRAD, IBM T.J. Watson Research Center
CHARLES A. ZRAKET, MITRE Corporation (retired)

NORMAN METZGER, Executive Director

Preface

In early 1994, acting through the Defense Authorization Act for FY 1994 (Public Law 103-160), Congress asked the National Research Council (NRC) to examine the status of the High Performance Computing and Communications Initiative (HPCCI). Broad-based interest in and support for the HPCCI exist. Given its scope and size, concerns had been raised about its goals, management, and progress. Congress asked that at a minimum the study address:

- The basic underlying rationale(s) for the program, including the appropriate balance between federal efforts and private-sector efforts;
- The appropriateness of its goals and directions;
- The balance between various elements of the program;
- The effectiveness of the mechanisms for obtaining the views of industry and users for the planning and implementation of the program;
- The likelihood that the various goals of the program will be achieved;
- The management and coordination of the program; and
- The relationship of the program to other federal support of high-performance computing and communications, including acquisition of high-performance computers by federal departments and agencies in support of the mission needs of these departments and agencies.

For this study the NRC's Computer Science and Telecommunications Board (CSTB) convened a committee of 12 members, expert on pioneering applications of computers and communications and the major components of the HPCCI: High-Performance Computing Systems, Advanced Software Technology and Algorithms, the National Research and Education Network, Basic Research and Human Resources, and Information Infrastructure Technology and Applications. Congress asked the committee to accelerate the normal NRC study process in order to provide an interim report by July 1, 1994, and a final report by February 1, 1995. The committee was able to meet this rapid turnaround by drawing on the knowledge and experience of its members as expressed in committee deliberations and by obtaining input from numerous outside experts.

The full committee met six times between March 10, 1994, and December 20, 1994, to hear more than 25 high-performance computing and communications users, builders, and scientists; to discuss the HPCCI in detail; and to produce this report. Additionally, smaller groups of committee members made site visits to discuss first-hand the use of high-performance technologies. These visits involved another 6 individuals from the Ford Motor Company and approximately 50 high-performance computing and communications users who had gathered at a workshop to discuss the use of high-performance systems in environmental research and simulation. In addition to examining the current status of the program, the committee considered the evolution of the HPCCI

and its goals and alternate government investment strategies related to technological development. The committee took into account varying perspectives on the initiative's goals and available assessments of progress toward achieving them.

The committee's interim report provided technical background and perspective on the overall development of high-performance computing and communications systems, as well as on the HPCCI, formally started in 1991.¹ The interim report made two recommendations: (1) strengthen the National Coordination Office to help it meet increasing future demands for program coordination and information functions; and (2) immediately appoint the congressionally mandated HPCCI Advisory Committee to provide broad-based, active input to the initiative. As this final report goes to press, both recommendations have yet to be acted on and thus require executive and legislative attention.

This final report, *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*, purposely adopted a broad perspective so as to examine the HPCCI within the context of the evolving information infrastructure and national economic competitiveness generally. Committee deliberations consistently pointed out the important contributions that computing and communications research have made to the nation's economy, scientific research, national defense, and social fabric. That research has nourished U.S. leadership in information technology goods, services, and applications. Traditionally, the most powerful computers and the fastest networks made many of those contributions. Recently, however, the widespread availability of significant computing and communications capabilities on the nation's desktops and factory floors has also produced many benefits. The broadening and interconnection of more and more computer-based systems call attention to research needs associated with system scale as well as performance.

The Committee to Study High Performance Computing and Communications: Status of a Major Initiative is grateful for the help, encouragement, and hard work of the NRC staff working with us: Marjory Blumenthal, Jim Mallory, Susan Maurizi, and Leslie Wade. Our meetings went smoothly because of their careful preparation. This report came together because of their attention and diligence. They patiently assembled sometimes conflicting text from diverse authors and helped reconcile it with the critique of our reviewers. They searched out facts of importance to our deliberations. Their excellent staff preparation helped us focus on the substance of our task.

Many others also made valuable contributions to the committee. In addition to individuals listed in Appendix F who briefed the committee, the committee appreciates inputs from Sally Howe and Don Austin (National Coordination Office); Bob Borchers, Paul Young, Dick Kaplan, and Robert Voight (National Science Foundation); Eric Cooper (FORE Systems); Stephen Squires (Advanced Research Projects Agency); Sandy MacDonald (National Oceanic and Atmospheric Administration); Robert Bonometti (Office of Science and Technology Policy); and Al Rosenheck (former congressional staffer). It is also grateful to the anonymous reviewers who helped to sharpen and focus the report with their insightful comments. Responsibility for the report, of course, remains with the committee.

¹Computer Science and Telecommunications Board (CSTB), National Research Council. 1994. *Interim Report on the Status of the High Performance Computing and Communications Initiative*. Computer Science and Telecommunications Board, Washington, D.C.

Finally, we want to acknowledge the contributions to our present task of research projects 30 years past. Of course this text was all word processed. Of course the charts were drafted on computers. Of course we used Internet communication nationwide to plan our meetings, share our thoughts, reconcile our differences, and assemble our report. We plugged our portable computers into a local network at each of our meetings, sending drafts to local laser printers. In short, we have partaken fully of the fruits of the HPCCI's precursors. We thank the visionaries of the past for our tools.

Frederick P. Brooks

Ivan E. Sutherland

Co-chairs

Committee to Study High Performance Computing and Communications:

Status of a Major Initiative

Contents

EXECUTIVE SUMMARY	1
1 U.S. LEADERSHIP IN INFORMATION TECHNOLOGY	13
Information Technology Is Central to Our Society, 13	
Information Technology Advances Rapidly, 14	
Retaining Leadership in Information Technology Is Vital to the Nation, 15	
The Federal Investment in Computing Research Has Paid Rich Dividends, 16	
Continued Federal Investment Is Necessary to Sustain Our Lead, 23	
Today the HPCCI Is the Umbrella for Most Government-sponsored Computing and Communications Research, 25	
Notes, 27	
2 THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS INITIATIVE	28
HPCCI: Goals and Emphases, 28	
Basic Objectives, 28	
Expanded Objectives, 30	
HPCCI Accomplishments, 31	
The Issue of Measurement, 31	
Better Computing and Computational Infrastructure, 32	
Increasing Researcher-Developer-User Synergy, 32	
Impact on Mission Agencies, 36	
Five Gigabit Testbed Projects: Collaboration and Impact, 36	
Evolution of HPCCI Goals and Objectives, 37	
Improving the Information Infrastructure, 37	
Evolving Research Directions and Relevance for the Information Infrastructure, 39	
Overall Computing and Communications R&D Planning, 41	
Toward a Better Balance, 42	
Moving Forward—Basic Issues, 42	
Balance of Private and Public Investment, 42	
Coordination Versus Management, 43	
Coordinating Structure, 45	
Budget, 49	
Notes, 49	

3 RECOMMENDATIONS	53
General Recommendations,	53
Recommendations on High-Performance Computing,	54
Recommendations on Networking and Information Infrastructure,	57
Recommendations on the Supercomputer Centers and Grand Challenge Program,	59
Recommendations on Coordination and Program Management in the HPCCI,	61
Comments Relating This Report's Recommendations for High-Performance Computing and Communications Research to Administration Priorities,	64
Notes,	66
BIBLIOGRAPHY	67
<i>APPENDIXES</i>	
A The High Performance Computing and Communications Initiative: Background,	77
B High-Performance Communications Technology and Infrastructure,	92
C Review of the High Performance Computing and Communications Initiative Budget,	97
D Current High Performance Computing and Communications Initiative Grand Challenge Activities,	105
E Accomplishments of National Science Foundation Supercomputer Centers,	108
F Individuals Providing Briefings to the Committee,	118

Industry, if left to itself, will naturally find its way to the most useful and profitable employment. Whence it is inferred that manufacturers, without the aid of government, will grow up as soon and as fast as the natural state of things and the interest of the community may require.

Against the solidity of this hypothesis . . . very cogent reasons may be offered . . . [including] the strong influence of habit; the spirit of imitation; the fear of want of success in untried enterprises; the intrinsic difficulties incident to first essays towards [competition with established foreign players]; the bounties, premiums, and other artificial encouragements with which foreign nations second the exertions of their own citizens. . . . To produce the desirable changes as early as may be expedient may therefore require the incitement and patronage of government.

—Alexander Hamilton, 1791, *Report on Manufactures*

Executive Summary

Information technology drives many of today's innovations and offers still greater potential for further innovation in the next decade. It is also the basis for a domestic industry of about \$500 billion,¹ an industry that is critical to our nation's international competitiveness. Our domestic information technology industry is thriving now, based to a large extent on an extraordinary 50-year track record of public research funded by the federal government, creating the ideas and people that have let industry flourish. This record shows that for a dozen major innovations, 10 to 15 years have passed between research and commercial application (see Figure ES.1). Despite many efforts, commercialization has seldom been achieved more quickly.

Publicly funded research in information technology will continue to create important new technologies and industries, some of them unimagined today, and the process will continue to take 10 to 15 years. Without such research there will still be innovation, but the quantity and range of new ideas for U.S. industry to draw from will be greatly diminished. Public research, which creates new opportunities for private industry to use, should not be confused with industrial policy, which chooses firms or industries to support. Industry, with its focus mostly on the near term, cannot take the place of government in supporting the research that will lead to the next decade's advances.

The High Performance Computing and Communications Initiative (HPCCI) is the main vehicle for public research in information technology today and the subject of this report. By the early 1980s, several federal agencies had developed independent programs to advance many of the objectives of what was to become the HPCCI. The program received added impetus and more formal status when Congress passed the High Performance Computing Act of 1991 (Public Law 102-194) authorizing a 5-year program in high-performance computing and communications. The initiative began with a focus on high-speed parallel computing and networking and is now evolving to meet the needs of the nation for widespread use on a large scale as well as for high speed in computation and communications. To advance the nation's information infrastructure there is much that needs to be discovered or invented, because a useful "information highway" is much more than wires to every house.

As a prelude to examining the current status of the HPCCI, this report first describes the rationale for the initiative as an engine of U.S. leadership in information technology and outlines the contributions of ongoing publicly funded research to past and current progress in developing computing and communications technologies (Chapter 1). It then describes and evaluates the HPCCI's goals, accomplishments, management, and planning (Chapter 2). Finally, it makes recommendations aimed at ensuring continuing U.S. leadership in information technology through wise evolution and use of the HPCCI as an important lever (Chapter 3). Appendixes A through F of the report provide additional details on and documentation for points made in the main text.

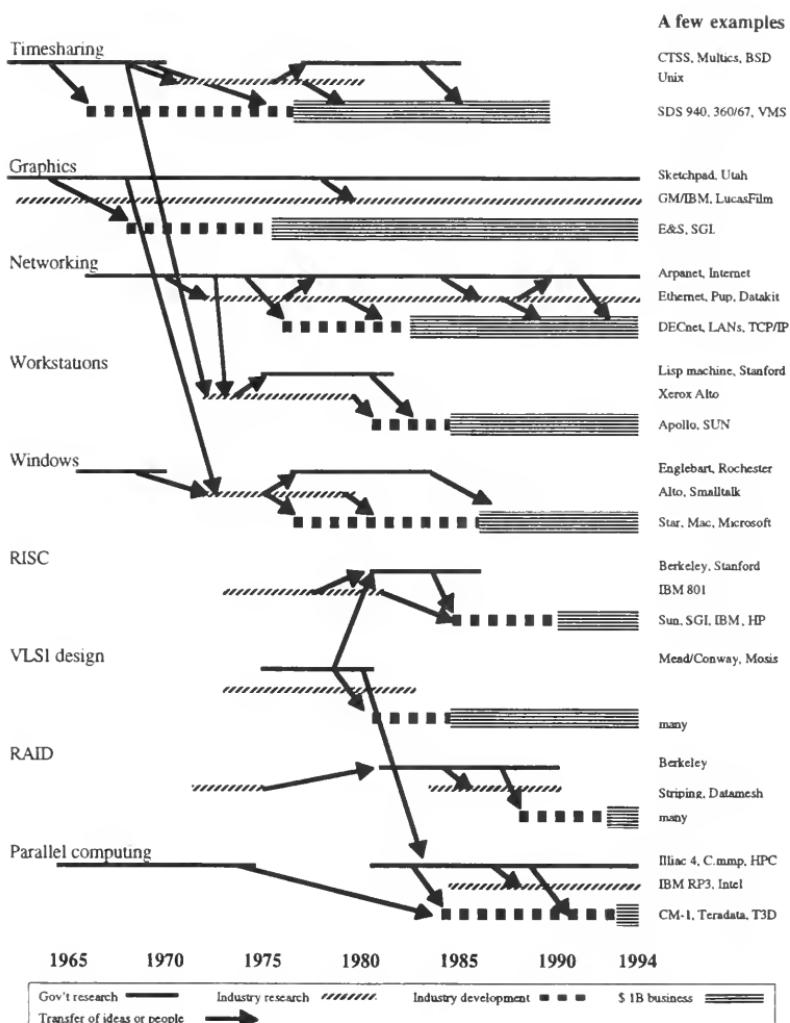


FIGURE ES.1 Government-sponsored computing research and development stimulates creation of innovative ideas and industries. Dates apply to horizontal bars, but not to arrows showing transfer of ideas and people.

INFORMATION TECHNOLOGY—FUNDAMENTAL FOR SOCIETY AND THE ECONOMY NOW AND TOMORROW

Computers, the devices that process information, affect our lives both directly and indirectly. Today, more than 70 million microcomputers are installed in the United States, and between one-fifth and one-third of U.S. households have one.² Entertainment, education, communications, medicine, government, and finance are using computers in more ways to enhance our lives directly through the provision of such services as distributed learning and remote banking. Computers are also used to make essential products and activities cheaper and better: airplanes, molded plastics, automobiles, medical imaging, and oil exploration are only a few of many examples. A broader benefit is the \$500 billion industry's creation of jobs, taxes, profits, and exports.

Clearly, the uses and applications of information technology will continue to grow. In fact, the information revolution has only just begun. Computers will become increasingly valuable to industries and to citizens as their power is tapped to recognize and simulate speech, generate realistic images, provide accurate models of the physical world, build huge automated libraries, control robots, and help with a myriad of other tasks. To do these things well will require both computing and communications systems many times more powerful than we have today. Ongoing advances in knowledge will constitute the foundation for building the systems and developing the applications that will continue to advance our quality of life and ensure strong U.S. leadership in information technology. Strong leadership in information technology in turn supports other sectors including industry, health, education, and defense by serving their needs for equipment, software, and know-how.

The Basis for Continuing Strength— A Successful Government-Industry Partnership

Federal investment in information technology research has played a key role in the U.S. capability to maintain its international lead in information technology. Starting in World War II publicly funded research has helped to stock the nation's storehouse of trained people and innovative ideas. But our lead is fragile. Leadership can shift in a few product generations, and because a generation in the computing and communications industry is at most 2 years, our lead could disappear in less than a decade.

Since the early 1960s the U.S. government has invested broadly in computing research, creating new ideas and trained people. The result has been the development of important new technologies for time-sharing, networking, computer graphics, human-machine interfaces, and parallel computing, as well as major contributions to the design of very large scale integrated circuits, fast computers and disk systems, and workstations (see Figure ES.1; see also Chapter 1, Box 1.2 for details). Each of these is now a multibillion-dollar business. From these successes we can learn some important lessons:

- *Research has kept paying off over a long period.*
- *The payoff from research takes time.* As Figure ES.1 shows, at least 10 years, more often 15, elapse between initial research on a major new idea and commercial success. This is still true in spite of today's shorter product cycles.
- *Unexpected results are often the most important.* Electronic mail and the "windows" interface are only two examples; Box 1.2 in Chapter 1 outlines more.

- *Research stimulates communication and interaction.* Ideas flow back and forth between research programs and development efforts and between academia and industry.
- *Research trains people,* who start companies or form a pool of trained personnel that existing companies can draw on to enter new markets quickly.
- *Doing research involves taking risks.* Not all public research programs have succeeded or led to clear outcomes even after many years. But the record of accomplishments suggests that government investment in computing and communications research has been very productive.

Government Support of Research Is Crucial

The information technology industry improves its products faster than most others: for the last 40 years a dollar has bought hardware with twice as much computation, storage, and communication every 18 to 24 months, offering a 100-fold gain every decade (Patterson and Hennessy, 1994, p. 21). This rate will continue at least for the next decade (see Chapter 1, Figure 1.1). Better hardware in turn makes it feasible to create software for new applications: electronic and mechanical design, climate mapping, digital libraries, desktop publishing, video editing, and telemedicine are just a few examples. Such applications are often brought to market by new companies such as Microsoft and Sun Microsystems, both of which produce revenues of more than \$4 billion per year (Computer Select, 1994) and neither of which existed 15 years ago.

The information technology industry is characterized by great importance to the economy and society, rapid and continuing change, a 10- to 15-year cycle from major idea to commercial success, and successive waves of new companies. In this environment a broad program of publicly funded research is essential for two reasons:

- First, industrial efforts cannot replace government investment in basic computing and communications research. Few companies will invest for a payoff that is 10 years away, and even a company that does make a discovery may postpone using it. The vitality of the information technology industry depends heavily on new companies, but new companies cannot easily afford to do research; furthermore, industry in general is doing less research now than in the recent past (Geppert, 1994; Corcoran, 1994). But because today's sales are based on yesterday's research, investment in innovation must go forward so that the nation's information industry can continue to thrive.
- Second, it is hard to predict which new ideas and approaches will succeed. The exact course of exploratory research cannot be planned in advance, and its progress cannot be measured precisely in the short term. The purpose of publicly funded research is to advance knowledge and create new opportunities that industry can exploit in the medium and long term, not to determine how the market should develop.

THE HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS INITIATIVE

Goals and Emphases

The HPCCI is the current manifestation of the continuing government research program in information technology, an investment that has been ongoing for more than 50 years. Although it emphasizes research in high-performance computing and communications, the HPCCI now has in its budget nearly all of the federal funding for computing research of any kind. The wisdom of this arrangement is doubtful.

The HPCCI was initiated to serve several broad goals (NCO, 1993):

- Extend U.S. leadership in high-performance computing and networking;
- Disseminate new technologies to serve the economy, national security, education, health care, and the environment; and
- Spur gains in U.S. productivity and industrial competitiveness.

The original plans to achieve these goals called for creating dramatically faster computers and networks, stretching their limits with Grand Challenge problems in scientific computing, setting up supercomputer centers with the machines and experts needed to attack these challenges, and training people to build and exploit the new technology. More recently the focus has been shifting toward broader uses of computing and communications.

High Performance

"High performance"—which involves bringing more powerful computing and communications technology to bear on a problem—has enabled advances on several fronts. High-performance systems, for example, deliver answers sooner for complex problems that need large amounts of computing. Timely and accurate forecasting of weather, mapping of oil reservoirs, and imaging of tumors are among the benefits encompassed by the goals listed above. But "high performance," which is broader than supercomputing, is a moving target because of the steady and rapid gains in the performance/cost ratio. Yesterday's supercomputer is today's personal computer; today's leading-edge communications technology will be among tomorrow's mainstream capabilities.

Information technology evolves as new and valuable applications are found for hardware that gets steadily more powerful and cheaper. To benefit, users need affordable hardware, but they also need the software that implements the new applications. Yet learning how to build software takes many years of experimentation. If this process starts only when the hardware has already become cheap, the benefits to users will be delayed by years. Research needs to treat today's expensive equipment as a time machine, learning how it will be used when it is cheap and widely available, as it surely will be tomorrow. Knowing how to use computers for new tasks sooner can help many industries to become more competitive.

To date, the HPCCI's focus has been mainly on speed, but speed is not the only measure of high performance. Both speed, measured today in billions of operations per second or billions of bits per second, and scale, measured by the number of millions of users served, are important research issues. However, for the nation's information infrastructure, scale now seems more difficult to achieve. Information technology can be thought of as a tent, with the height of the center pole as speed and the breadth of the base as scale. Widening the tent to allow more work on scale without decreasing the work on speed requires more cloth; with the same resources, widening

the tent would sacrifice research on speed for research on scale. This report recommends ways to reallocate funds within the HPCCI so as to accommodate greater emphasis on scale.

Accomplishments to Date

The HPCCI has focused mainly on parallel computing, networking, and development of human resources. Building on progress in research begun before the HPCCI, work and accomplishments to date reveal two key trends: better computing and computational infrastructure and increasing researcher-developer-user synergy.

Despite the difficulty of measuring impact at this early stage, it is the committee's judgment that the HPCCI has been generally successful so far. That assessment is necessarily qualitative and experiential now. Because the HPCCI is only 3 years old, results that can be measured in dollars should not be expected before the next decade.

The HPCCI has contributed substantially to the development, deployment, and understanding of computing and communications facilities and capabilities as infrastructure. It has helped transform understanding of how to share resources and information, generating proofs of concept and understanding that are of value not only to the scientific research community but also to the economy and society at large.

In parallel computing the fundamental challenge is not building the machines, but learning how to program them. Pioneering users and their software developers must be motivated by machines that are good enough to reward success with significant speedups.³ For this reason, a great deal of money and effort have had to be spent to obtain parallel machines with the potential to run much faster than existing supercomputers. From the base built by the HPCCI, much has been learned about parallel computing.

The HPCCI has fostered productive interactions among the researchers and developers involved in creating high-performance computing and communications technology and researchers who use the technology. Building on the varying perspectives of the three groups, complex problems are being solved in unique ways. In particular, the HPCCI has funded cross-disciplinary teams associated with the Grand Challenge projects to solve complex computational problems and produce essential new software for the new parallel systems.

More specifically, the HPCCI has:

- Increased the nation's stock of expertise by educating new students and attracting new researchers;
- Made parallel computing widely accepted as the practical route to achieving high-performance computing;
- Demonstrated the feasibility of and initiated deployment of parallel databases;
- Driven progress on Grand Challenge problems in disciplines such as cosmology, molecular biology, chemistry, and materials science. Parallel computation has enhanced the ability to attack problems of great complexity in science and engineering;
- Developed new modes of analyzing and visualizing complex data sets in the earth sciences, medicine, molecular biology, and engineering, including creating virtual reality technologies. Many supercomputer graphic techniques of the 1980s are now available on desktop graphics workstations;

- Through the gigabit network testbeds associated with the National Research and Education Network component, demonstrated the intimate link between computing and communications systems;
- Built advanced networks that are the backbone of the Internet and the prototypes for its further evolution into the basis for a broader information infrastructure;
- Deployed a high-speed backbone that has kept up with the yearly doubling of the size of the Internet, and organized the impending transition of this backbone away from government funding; and
- Created the Mosaic browser for the World Wide Web, the first new major application in many years that promises to greatly increase access to the resources available on the Internet. This was an entirely unexpected result.

Evolution

A large-scale, integrated information infrastructure designed to serve the entire nation is becoming a high priority for government and industry as well as a source of challenges for research. Complex systems with millions of users pose many problems: performance, management, security, interoperability, compatible evolution of components, mobility, and reliability are only a few. Today's technology can solve these problems for systems with a few thousand users at most; to do so for millions or hundreds of millions of users is far beyond the current state of the art.⁴ Providing users with high-bandwidth connections is itself a major problem, but it is only the beginning. There is a wide gap between enabling a connection and providing a rich array of useful and dependable services.

Because the HPCCI has become the rubric under which virtually all of the nation's research in information technology is conducted, it is not surprising that its focus has been changing in response to past successes, new opportunities, and evolving societal needs. The recently added Information Infrastructure Technology and Applications (IITA) program, broadly construed, addresses many of the problems just mentioned; it is already the largest component of the HPCCI,⁵ and its continued evolution should be encouraged.

But with the policy focus—in the government, the press, and in most of the agencies—centered on information infrastructure,⁶ high-performance computing seems to have been downplayed. The committee emphasizes the importance of retaining the HPCCI's momentum at just the time when its potential to support improvement in the nation's information infrastructure is most needed.

Organization

Several federal agencies participate in the HPCCI, most notably the National Science Foundation (NSF), the Advanced Research Projects Agency (ARPA), the National Aeronautics and Space Administration (NASA), and the Department of Energy (DOE) (see Appendix A for a full list). Because of its successes the HPCCI has become a model for multiagency collaboration and for the "virtual agency" concept advanced through the *National Performance Review* (Gore, 1993). Each participating agency retains responsibility for its own part of the program, but the agencies work together in joint funding of projects, such as the federal portions of the Internet; joint reviews of grants and contracts, such as the NSF-ARPA-NASA digital library initiative; joint testbeds; and

consortia, such as the consortium for Advanced Modeling of Regional Air Quality that joins six federal agencies with several state and local governments.

The HPCCI supports a diverse set of contractors at universities, companies, and national laboratories throughout the country. It provides project funding in varying amounts through contracts, grants, and cooperative agreements awarded according to diverse methods. This diversity is healthy because it allows many views to compete, resulting in a broad research program that ensures a continuing flow of advances in information technology.

Some have argued for a more centrally managed program, with thorough planning, precise milestones, and presumably no wasted effort. Tighter management would cost more in bureaucracy and turf wars, but the essential question is whether it would produce better or worse results for the money spent. The committee believes that because of the long time scale of research, diversity is essential for success. No one person or organization is either smart or lucky enough to plan the best program, no single approach is best, and success often comes in unanticipated ways. Because it is a national research program and because of the many different but interdependent underlying technologies, the HPCCI is necessarily and properly far more diverse than a focused effort such as the Apollo moon landing program or a commercial product development program.

In contrast to central management, coordination enhances the benefits of diversity by helping to prevent unintended duplication, redundancy, and missed opportunities. The HPCCI's National Coordination Office (NCO) serves this purpose, aiding interagency cooperation and acting as liaison for the initiative to the Congress, other levels of government, universities, industry, and the public. Its efforts are reflected in its impressive FY 1994 and FY 1995 "Blue Books" describing the program's activities and spending.⁷ Strengthening the NCO and appointing an advisory committee, as recommended in the committee's interim report (CSTB, 1994c), would facilitate regular infusions of ideas and advice from industry and academia and enable better communication of the HPCCI's goals and accomplishments to its many constituents. This committee should consist of a group of recognized experts that is balanced between academia and industry and balanced with respect to application areas and the core technologies underlying the HPCCI.

Budget

Because it grew from earlier programs, a significant portion of the HPCCI budget is not new money. The budget grew from a base of \$490 million in preexisting funding in FY 1992 to the \$1.1 billion requested for FY 1995.⁸ Each year agencies have added to the base by moving budgets for existing programs into the HPCCI and by reprogramming existing funds to support the HPCCI. Congress has also added funding each year to start new activities or expand old ones.

The result is that much of the \$1.1 billion requested for FY 1995 is money that was already being spent on computing and communications in FY 1992. The request has three elements: (1) funds for activities that predate the HPCCI and were in the FY 1992 base budget, (2) funds for activities that have since been designated as part of the HPCCI, and (3) new funds for new activities or for growth. Although dissecting the budget in this way would shed light on the program, the committee was unable to do so because each participating agency treats the numbers differently.

It appears that the FY 1995 request breaks down roughly as one-third for applications, one-third to advance the essential underlying computing and communications technologies, one-quarter for computing and communications infrastructure, and small amounts for education and electronics (see Appendix C).

THE FUTURE OF THE HPCCI: RECOMMENDATIONS

The committee believes that strong public support for a broadly based research program in information technology is vital to maintaining U.S. leadership in information technology. Incorporating this view of the importance and success of the government's investment in research, the 13 recommendations that follow address five areas: general research program, high-performance computing, networking and information infrastructure, the supercomputer centers and the Grand Challenge projects, and program coordination and management. Within each area the recommendations are presented in priority order.

General Recommendations

1. Continue to support research in information technology. Ensure that the major funding agencies, especially the National Science Foundation and the Advanced Research Projects Agency, have strong programs for computing and communications research that are independent of any special initiatives.

The government investment in computing research has yielded significant returns. Ongoing investment, *at least as high as the current dollar level*, is critical both to U.S. leadership and to ongoing innovation in information technology. Today the HPCCI supports nearly all of this research, an arrangement that is both misleading and dangerous: misleading because much important computing research addresses areas other than high performance (even though it may legitimately fit under the new IITA component of the HPCCI), and dangerous because reduced funding for the HPCCI could cripple all of computing research. The "war on cancer" did not support all of biomedical research, and neither should the HPCCI or any future initiative on national infrastructure subsume all of computing research.

2. Continue the HPCCI, maintaining today's increased emphasis on the research challenges posed by the nation's evolving information infrastructure. The new Information Infrastructure Technology and Applications program of the HPCCI focuses on information infrastructure topics, which are also addressed in the initiative's other four components. The committee supports this continued evolution, which will lead to tangible returns on existing and future investments in basic hardware, networking, and software technologies.

High-Performance Computing

3. Continue funding a strong experimental research program in software and algorithms for parallel computing machines. Today a crucial obstacle to widespread use of parallel computing is the lack of advanced software and algorithms. Emphasis should be given to research on developing and building usable applications-oriented software systems for parallel computers. Avoid funding the transfer ("porting") of existing commercial applications to new parallel computing machines unless there is a specific research need.

4. Stop direct HPCCI funding for development of commercial hardware by computer vendors and for "industrial stimulus" purchases of hardware. Maintain HPCCI support for precompetitive research in computer architecture; this work should be done in universities or in university-industry collaborations and should be driven by the needs of system and application software. HPCCI funding for stimulus purchase of large-scale machines has been reduced, as has the funding of hardware development by vendors. The committee supports these changes, which should continue except when a mission need demands the development of nonstandard hardware.

Public research is best done in universities. Not only are academic organizations free to think about longer-term issues, but they also stimulate technology transfer through publication and placement of graduates. The national experience supports Vannevar Bush's basic tenet: publicly funded research carried out in universities produces a diversity of excellent ideas, trained people, research results, and technologies that can be commercially exploited (OSRD and Bush, 1945).

5. Treat development of a teraflop computer as a research direction rather than a destination. The goal of developing teraflop capability has served a valuable purpose in stimulating work on large-scale parallelism, but further investment in raw scalability is inappropriate except as a focus for precompetitive, academic research. Industrial development of parallel computers will balance the low cost of individual, mass-produced computing devices against the higher cost of communicating between them in a variety of interesting ways. In the near future a teraflop parallel machine will be built when some agencies' mission requirements correspond to a sufficiently economical commercial offering. Continued progress will surely lead to machines even larger than a teraflop.

Networking and Information Infrastructure

New ideas are needed to meet the new challenges underlying development of the nation's information infrastructure. The HPCCI can contribute most by focusing on the underlying research issues. This shift has already begun, and it should continue.

This evolution of the research agenda, which would support improvement of the nation's information infrastructure, is partly under way: in the *FY 1995 Implementation Plan* (NCO, 1994, p. 15), over one-quarter of the NSF and ARPA HPCCI funding is focused on the IITA component, and activities in other components have also evolved consistent with these concerns. The committee supports this increased emphasis.

6. Increase the HPCCI focus on communications and networking research, especially on the challenges inherent in scale and physical distribution. An integrated information infrastructure that fully serves the nation's needs cannot spring full-grown from what we already know. Much research is needed on difficult problems related to size, evolution, introduction of new systems, reliability, and interoperability. Much more is involved than simply deploying large numbers of boxes and wires. For example, both hardware and software systems must work efficiently to handle scheduling; bandwidth optimization for transmission of a range of data formats, including real-time audio and video data; protocol and format conversion; security; and many other requirements.

7. Develop a research program to address the research challenges underlying our ability to build very large, reliable, high-performance, distributed information systems based on the existing HPCCI foundation. Although a comprehensive vision of the research needed for advancing the nation's information infrastructure has not yet been developed, three key areas for research are scalability, physical distribution, and interoperative applications.

8. Ensure that research programs focusing on the National Challenges contribute to the development of information infrastructure technologies as well as to the development of new applications and paradigms. This dual emphasis contrasts with the narrower focus on scientific results that has driven work on the Grand Challenges.

Supercomputer Centers and Grand Challenge Program

The NSF supercomputer centers have played a major role in establishing parallel computing as a full partner with the prior paradigms of scalar and vector computing by providing access to

state-of-the-art computing facilities. NSF should continue to take a broad view of the centers' mission of providing access to high-performance computing and communications resources, including participating in research needed to improve software for parallel machines and to advance the nation's information infrastructure.

The committee recognizes that advanced computation is an important tool for scientists and engineers and that support for adequate computer access must be a part of the NSF research program in all disciplines. The committee did not consider the appropriate overall funding level for the centers. However, the committee believes that NSF should move to a model similar to that used by NASA and DOE for funding general access to computing. The committee prefers NASA's and DOE's approach to funding supercomputer centers, where HPCCI funds are used only to support the exploration and use of new computing architectures, while non-HPCCI funds are used to support general access.

9. The mission of the National Science Foundation supercomputer centers remains important, but the NSF should continue to evaluate new directions, alternative funding mechanisms, new administrative structures, and the overall program level of the centers. NSF could continue funding of the centers at the current level or alter that level, but it should continue using HPCCI funds to support applications that contribute to the evolution of the underlying computing and communications technologies, while support for general access by application scientists to maturing architectures should come increasingly from non-HPCCI funds.

10. The Grand Challenge program is an innovative approach to creating interdisciplinary and multi-institutional scientific research teams; however, continued use of HPCCI funds is appropriate only when the research contributes significantly to the development of new high-performance computing and communications hardware or software. Grand Challenge projects funded under the HPCCI should be evaluated on the basis of their contributions both to high-performance computing and communications technologies and to the application area. Completion of the Grand Challenge projects will provide valuable insights and demonstrate the capabilities of new high-performance architectures in some important applications. It will also foster better collaboration between computer scientists and computational scientists. The committee notes that a large share of HPCCI funding for the Grand Challenges currently comes from the scientific disciplines involved. However, the overall funding seems to come entirely from HPCCI-labeled funds. For the same reasons outlined in Recommendation 9, the committee sees this commingled support as unhealthy in the long run and urges a transition to greater reliance on scientific disciplinary funding using non-HPCCI funds.

Coordination and Program Management

11. Strengthen the HPCCI National Coordination Office while retaining the cooperative structure of the HPCCI and increasing the opportunity for external input. Immediately appoint the congressionally mandated advisory committee intended to provide broad-based, active input to the HPCCI, or provide an effective alternative. Appoint an individual to be a full-time coordinator, program spokesperson, and advocate for the HPCCI.

In making this recommendation, the committee strongly endorses the role of the current NCO as supporting the mission agencies rather than directing them. The committee believes it vital that the separate agencies retain direction of their HPCCI funds. The value of interagency cooperation outweighs any benefits that might be gained through more centralized management.

Diverse management systems for research should be welcomed, and micromanagement should be avoided. In the past, choosing good program officers and giving them freedom to operate independently have yielded good value, and the committee believes it will continue to do so.

Furthermore, independence will encourage diversity in the research program, thus increasing opportunities for unexpected discoveries, encouraging a broader attack on problems, and ensuring fewer missed opportunities.

12. Place projects in the HPCCI only if they match well to its objectives. Federal research funding agencies should promptly document the extent to which HPCCI funding is supporting important long-term research areas whose future funding should be independent of the future of the HPCCI.

A number of preexisting agency programs have entered the HPCCI, with two effects: the HPCCI's budget appears to grow faster than the real growth of investment in high-performance computing and communications research, and important programs such as basic research in computing within NSF and ARPA may be in jeopardy should the HPCCI end.

13. Base mission agency computer procurements on mission needs only, and encourage making equipment procurement decisions at the lowest practical management level. This recommendation applies equally to government agencies and to government contractors. It has generally been best for an agency to specify the results it wants and to leave the choice of specific equipment to the contractor or local laboratory management.

NOTES

1. See U.S. DOC (1994); the Department of Commerce utilizes data from the U.S. Bureau of the Census series, the *Annual Survey of Manufactures*. It places the value of shipments for the information technology industry at \$421 billion for 1993. This number omits revenue from equipment rentals, fees for after-sale service, and mark-ups in the product distribution channel. It also excludes office equipment in total. It includes computers, storage devices, terminals and peripherals; packaged software; computer program manufacturing, data processing, information services, facilities management, and other services; and telecommunications equipment and services.

See also CBEMA (1994); CBEMA values the worldwide 1993 revenue of the U.S. information technology industry at \$602 billion. In addition to including office equipment, it shows larger revenues for information technology hardware and telecommunications equipment than does the Department of Commerce.

2. Microcomputers (personal computers) are defined as computers with a list price of \$1,000 to \$14,999; see CBEMA (1994), pp. 60-61. Forrester Research Inc. (1994, pp. 2-3) estimates the share of households with PCs at about 20 percent, based on its survey of households and Bureau of Census data. Forrester's model accounts for retirements of older PCs and for households with multiple PCs. This is a lower estimate than the Software Publishing Association's widely cited 30 percent share. By definition, the microcomputer statistics exclude small computers and other general-purpose and specialized devices that also make use of microprocessors and would be counted in a more comprehensive measurement of information technology.

3. Earlier experience with three isolated computers, "Illiad 4" (built at the University of Illinois) and "C.mmp" and "C.m**" (both built at Carnegie Mellon University), bears out this point.

4. Of course, systems specialized for a single application or for homogenous technology, such as telephony, serve millions of users, but what is now envisioned is more complex and heterogeneous, involving integration of multiple services and systems.

5. The other four programs of the HPCCI are Advanced Software Technology and Algorithms, Basic Research and Human Resources, High-Performance Computing Systems, and the National Research and Education Network.

6. Notably, references to the computing portion of the HPCCI have been overshadowed recently by the ubiquity of speeches and documents devoted to the notion of a national information infrastructure (NII). The NII has also been featured in the titles of the 1994 and 1995 Blue Books.

7. Each year beginning in 1991 the director of the Office of Science and Technology Policy submits a report on the HPCCI to accompany the president's budget. The FY 1992, FY 1993, and FY 1994 books were produced by the now-defunct Federal Coordinating Council for Science, Engineering, and Technology; the FY 1995 report was produced by the NCO (acting for the Committee on Information and Communications). The report describes prior accomplishments and the future funding and activities for the coming fiscal year. These reports have collectively become known as "Blue Books" after the color of their cover.

8. NCO (1994), p. 15. Note that figures represent the President's requested budget authority for FY 1995. Actual appropriated levels were not available at press time. Because the HPCCI is synthesized as a cross-cutting multiagency initiative, there is no separate and identifiable "HPCCI appropriation."

1

U.S. Leadership in Information Technology

Information technology is central to our economy and society. The United States has held a commanding lead in this arena, a lead that we must maintain. Meeting the challenges posed by rapid, worldwide change will continue to require our best efforts to advance the state of the art in computing and communications technology. Now, as in the past, our ability to lead requires an ongoing strong program of long-term research. The federal government has supported such research for 50 years with great success.

Today, the High Performance Computing and Communications Initiative (HPCCI) is the multiagency cooperative effort under which most of this research is funded. For this reason, any discussion of the HPCCI must be grounded in an understanding of the role and nature of information technology, the information industry, and the nation's research program in computing. These issues are the subject of this first chapter of the report.

INFORMATION TECHNOLOGY IS CENTRAL TO OUR SOCIETY

Computers affect our lives enormously. We use them directly for everyday tasks such as making an airline reservation, getting money from an automated teller machine, or writing a report on a word processor. We also enjoy many products and services that would not be possible without digital computing and communications.

The direct use of computing is growing rapidly. Personal computers are already pervasive in our economy and society: in the United States over 70 million are installed, and between one-fifth and one-third of U.S. households have a computer.¹ The increasing popularity of personal computing is but the tip of the iceberg; education, communication, medicine, government, manufacturing, transportation, science, engineering, finance, and entertainment all increasingly use digital computing and communications to enhance our lives directly by offering improved goods and services.

Indirectly, computing and communications are used to make many products cheaper and better. Without computers connected by communication networks, designing the newest U.S. jetliner, the Boeing 777, within acceptable cost and time constraints would have been impossible. Advanced hardware and advanced software working together substituted computer models for expensive and time-consuming physical modeling. The design of complex plastic molded parts, now routinely used for many products, depends on computer simulation of plastic flow and computer control of die making. Automobile engines rely on embedded computers for fuel economy and emission control, and doctors use computer-assisted tomography (CAT) scanners to

see inside the body. Computers help us understand and tap Earth's resources: our oil is found by computer analysis of geologic data. Interconnected computer systems underlie our entire financial system, enabling electronic funds transfer and services such as home banking. Digital communication extends to a large and rapidly increasing number of businesses, educational institutions, government agencies, and homes.

Originally devices for computation and business data processing, computers now are tools for information access and processing in the broadest sense. As such, they have become fundamental to the operation of our society, and computing and communications have come to be labeled widely as "information processing."

Remarkably, given its already enormous direct and indirect impact, information technology is being deployed in our society more rapidly now than at any time in the past.² If this momentum is sustained, then digital technology and digital information—the digital information revolution—will offer a huge range of new applications, create markets for a wide variety of new products and services, and yield a broad spectrum of benefits to all Americans.

INFORMATION TECHNOLOGY ADVANCES RAPIDLY

The information industry improves its products with amazing speed. For several decades—powered by federal and industrial research and development (R&D) investments in computer science, computer engineering, electrical engineering, and semiconductor physics—each dollar spent on computation, storage, and communication has bought twice the performance every 18 to 24 months. Over the course of each decade, this sustained rate of progress results in a 100-fold improvement, as Figure 1.1 shows for processor speed and disk storage capacity. With continued investment, we can sustain this rate of progress for at least the next decade. Such rapid improvement is possible because of the nature of information and of the technologies required to process it: integrated circuits, storage devices, and communications systems (Box 1.1). Significant improvements in hardware performance in turn make it feasible to create the software required for computers to do new things—electronic and mechanical design, desktop publishing, video editing, modeling of financial markets, creation of digital libraries, and the practice of telemedicine, for example.

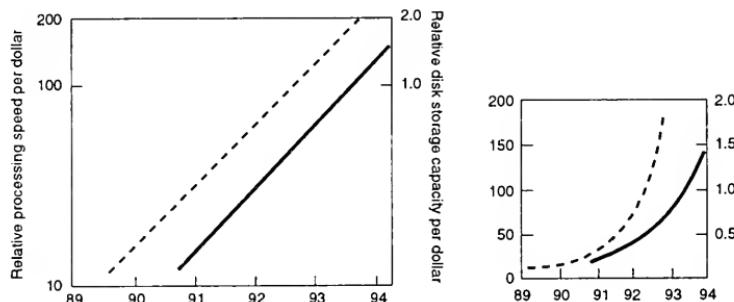


FIGURE 1.1 Increase in performance per dollar of processor speed and disk storage from 1989 to 1994, shown on a semilog scale. (The right-hand graph uses a linear scale to emphasize the compound effect of successive doublings.)

**BOX 1.1 What Drives the Progress
in Information Technology?**

- *Integrated circuits improve rapidly.* Computers are made from integrated circuits, each component of which gets half as high and half as wide every 5 years, with the result that the same area can hold four times as many components at the same cost. Also, each component runs twice as fast, and the circuit chips get bigger. The physical limits to this progress are still far off.
- *New designs take advantage of advances in integrated circuits.* Today's microprocessors and memories are made from very large scale integrated (VLSI) circuits. Although modern VLSI microprocessors and memories may have 10 million components, they are actually designed in no more time than the integrated circuits of a decade ago that had only 100,000 components. Advances in designs and design tools—in our ability to master complexity—have been, and will continue to be, essential to taking advantage of advances in integrated circuit technology.
- *It is cheap to make more devices, and the same integrated circuit foundries can make many different devices.* The marginal cost of building more computers is small, because the cost of raw materials is low and the components are mass-produced. Further, although an integrated circuit foundry may be expensive to build, it can make many different products, just as a printing press can print many different books. Because the same process is used over and over again, improvements in this process have enormous effects on product cost and quality.
- *New designs can quickly become products.* A new digital system is usually built in an existing foundry that operates directly from a digital description of the design. Investing in prototypes is not necessary, because it is possible to simulate a product accurately and automatically from the design.
- *Dramatic system advances enable dramatic application advances.* The fact that computing power per dollar doubles every 18 months means that capabilities can migrate from the high end to the consumer rapidly and predictably. It is R&D investment at the high end that creates these capabilities. Today's desktop workstation was the supercomputer of a mere decade ago. Today's solutions to the problems of large-scale parallelism will enable us to solve tomorrow's mass-market problems of on-chip parallelism.

Rapid progress has produced successive waves of new companies in diverse areas related to information technology and its applications: integrated circuits, computer hardware, computer software, communications, embedded systems, robotics, video on demand, and others. Many of today's major hardware and software firms, including Sun Microsystems and Microsoft, both with revenues of more than \$4 billion per year (Computer Select, 1994), did not exist 15 years ago, and none except IBM and Unisys were in the information business 40 years ago. Many of these new companies have drawn on ideas and people from federally funded research projects.

**RETAINING LEADERSHIP IN INFORMATION TECHNOLOGY
IS VITAL TO THE NATION**

The U.S. lead in information technology has brought benefits that are clearly valuable to the nation:

- *The jobs and profits generated by the industry itself.* The information technology industry is big: Revenues attributable to hardware and software sales plus associated services were on the order of \$500 billion in 1993.³ Due to the limitations of what is actually counted in

any given estimate, this figure may be conservative. The jobs and profits that the information technology industry complex delivers are worth careful preservation.

- *The advantages that other sectors gain from early access to the best information technology, ahead of our international competitors.* Learning how to use computers for new tasks sooner allows firms to capture new market segments and compete more efficiently in old ones. U.S. competitiveness in engineering, manufacturing, transportation, financial services, and a host of other areas depends on U.S. competitiveness in information technology (CSTB, 1994b, 1995).

- *The benefits that our citizens gain from information technology in their daily lives.* Benefits are evident in education, medicine, finance, communication, entertainment, information services, and other areas. The lives of Americans are improved by 24-hour banking, improved fuel economy in automobiles, noninvasive medical diagnosis, and hundreds of uses of computers for generating film sequences and as the basis for computerized games. The reach and impact of such applications are increasing rapidly.

Our lead in information technology is fragile, and it will slip away if we fail to adapt. Leadership has often shifted in a few product generations, and a generation in the information industry can be less than 2 years. As a nation we must continue to foster the flow of fresh ideas and trained minds that have enabled the U.S. information technology enterprise as a whole to remain strong despite the fate of individual companies. International competition is fierce, and it is likely to increase. In the mid-1980s Japan and Europe recognized the strategic importance of information technology and began investing massively in the Fifth Generation and Esprit efforts, respectively. Although these efforts did not yield new technologies to rival our own, their very significant investments in research and technology infrastructure are laying the foundation for future challenges.⁴

Today, through the efforts of government, academia, and industry, our nation retains its lead and continues to enjoy the benefits associated with it. Although many factors contributed, the committee believes that federal investment in the advancement of information technology has played a key role.

Indeed, for nearly 50 years federal investment has helped to train the people and stimulate the ideas that have made today's computers and many of their applications possible. Federal support early in the life cycle of many ideas has advanced them from novelties to convincing demonstrations that attract private investment to products and services that ultimately add to the quality of U.S. life.

THE FEDERAL INVESTMENT IN COMPUTING RESEARCH HAS PAID RICH DIVIDENDS

In the 1940s and 1950s, much of the federal investment in computing research was for defense (Flamm, 1988). Technical needs such as fire control and intelligence needs such as cryptography and mission planning required great computing power.

Since the early 1960s the federal government has invested more broadly in computing research, and these investments have profoundly affected how computers are made and used, contributed to the development of innovative ideas and training of key people, and led to the kinds of advances sampled in Table 1.1. Figure 1.2 shows the corresponding timelines for progress from research topics commercially successful applications. Notable among the government efforts stimulating many of these advances have been the Advanced Research Projects Agency's (ARPA's)

VLSI program, an initiative of the past decade (Box 1.2), and federal sponsorship of research in laying the groundwork for the now-ubiquitous field of computer graphics (Box 1.3) and for sophisticated database systems (Box 1.4).

TABLE 1.1 Some Successes of Government Funding of Computing and Communications Research

Topic	Goal	Unanticipated Results	Today
Timesharing	Let many people use a computer, each as if it were his or her own, sharing the cost.	Because many people kept their work in one computer, they could easily share information. Reduced cost increased the diversity of users and applications.	Even personal computers are timeshared among multiple applications. Information sharing is ubiquitous; shared information lives on "servers."
Computer networking	Load-sharing among a modest number of major computers	Electronic mail; widespread sharing of software and data; local area networks (the original networks were wide-area); the interconnection of literally millions of computers	Networking has enabled worldwide communication and sharing, access to expertise wherever it exists, and commerce at our fingertips.
Workstations	Enough computing to make interactive graphic useful	Displaced most other forms of computing and terminals; led directly to personal computers and multimedia	Millions in use for science, engineering, and finance
Computer graphics	Make pictures on a computer.	"What you see is what you get" and hypermedia documents	Almost any image is possible. Realistic moving images made on computers are routinely seen on television and were used effectively in the design of the Boeing 777.
"Windows and mouse" user interface technology	Easy access to many applications and documents at once	Dramatic improvements in overall ease of use; the integration of applications (e.g., spreadsheets, word processors, and presentation graphics)	The standard way to use all computers
Very large integrated circuit design	New design methods to keep pace with integrated circuit technology	Easy access to "silicon foundries"; a renaissance in computer design	Many more schools training VLSI designers; many more companies using this technology
Reduced Instruction Set Computers (RISC)	Computers 2 to 3 times faster	Dramatic progress in the "co-design" of hardware and software, leading to significantly greater performance	Millions in use; penetration continues to increase
Redundant Arrays of Inexpensive Disks (RAID)	Faster, more reliable disk systems	RAID is more economical as well: massive data repositories ride the price/performance wave of personal computers and workstations.	Entering the mainstream for large-scale data storage; will see widespread commercial use in digital video servers

continues

TABLE 1.1—continued

Topic	Goal	Unanticipated Results	Today
Parallel computing	Significantly faster computing to address complex problems	Parallel desktop server system; unanticipated applications such as transaction processing, financial modeling, database mining, and knowledge discovery in data	Many computer manufacturers include parallel computing as a standard offering.
Digital libraries	Universal, multimedia (text, image, audio, video) access to all the information in large libraries; an essential need is tools for discovering and locating information	Pending development	Beginning development

From this record of success we can draw some important conclusions:

- *Research pays off for an extended period.* The federal investment and the payoff, including the spawning of numerous corporations and multibillion-dollar industries, have been continuous for decades.
- *Unanticipated results are often the most important results.* Information sharing is an unanticipated result of timesharing; what-you-see-is-what-you-get displays and hypermedia documents are unanticipated results of computer graphics; electronic mail is an unanticipated result of networking.
- *The fusion of ideas multiplies their effect.* Distributed systems, such as automated teller machine networks, combine elements of timesharing, networking, workstations, and computer graphics. Personal digital assistants, the emerging generation of truly portable computers, combine these elements with new networking and power-management technologies.
- *Research trains people.* A major output of publicly supported research programs has been people. Some develop or create a new concept and start companies to commercialize their knowledge. Others form a pool of expertise that allows new or existing companies to move quickly into new technologies.
- *The synergy among industry, academia, and government has been highly effective.* The flow of ideas and people between government-sponsored and commercial programs is suggested in Figure 1.2.

**BOX 1.2 An Example of a Successful
Federal R&D Program: The ARPA VLSI Program**

The ARPA VLSI program began in the late 1970s. This program, inspired by the ground-breaking work of Carver Mead and Lynn Conway, envisioned that integrated circuit technology could be made available to system designers and that this would have tremendous impact. The program funded academic research activities as well as the Metal Oxide Semiconductor Implementation Service (MOSIS). MOSIS provided low-cost, fast-turnaround VLSI fabrication services to the research community; established by ARPA, it was expanded and had access broadened by the National Science Foundation. The ARPA VLSI program is widely regarded to have been a tremendous success. Among its notable achievements are the following:

- Developed the concept of the multichip wafer, which allowed multiple designs to share a single silicon fabrication run. Together with tools developed to assemble the designs and provide services for digital submission of chip designs, this capability made the concept of a low-cost, fast-turnaround silicon foundry a reality. Several companies were formed based on these ideas, with VLSI Technology Inc. being the best known.
- Stimulated development of the Geometry Engine and Pixel Planes projects, which used the capabilities of VLSI to create new capabilities in low-cost, high-performance three-dimensional graphics. The Geometry Engine project formed the basis of Silicon Graphics Inc. Pixel planes technology is licensed to Irix and Division.
- Stimulated development of Berkeley UNIX, which was funded to provide a research platform for the VLSI design tools. This version of UNIX eventually became the basis for the operating system of choice in workstations, servers, and multiprocessors. UNIX went on to become the most widely used vendor-independent operating system, with the code developed at Berkeley being key to this development.
- Accelerated understanding of the importance of low-cost, high-quality graphics for VLSI design, inspiring the creation of the Stanford University Network (SUNI) workstation project. Together with the UNIX development from Berkeley, this technology formed the basis for Sun Microsystems.
- Developed two of the three RISC experiments, the Berkeley RISC project and the Stanford MIPS project, which were major parts of the VLSI program inspired by the possibilities of VLSI technology. These technologies formed the basis for many RISC designs, including those of MIPS Computer Systems (now owned by Silicon Graphics Inc.) and Sun Microsystems.
- Sponsored extensive developments in computer-aided design (CAD) tool design. These led to revolutionary improvements in CAD technology for layout, design rule checking, and simulation. The tools developed in this program were used extensively in both academic research programs and in industry. The ideas were developed in commercial implementations by companies such as VLSI Technology, Cadnetix, and more recently, Synopsis.

Overall, the ARPA VLSI program was a landmark success, not only in creating new technologies and revolutionizing the computer industry, but also in forming the basis for major new industrial technologies and a number of companies that have become major corporations.*

*Interestingly, the success of the ARPA VLSI program stands in sharp contrast to the Department of Defense VHSC program, which based entirely in industry and is generally regarded to have had only modest impact either in developing new technologies or in accelerating technology.

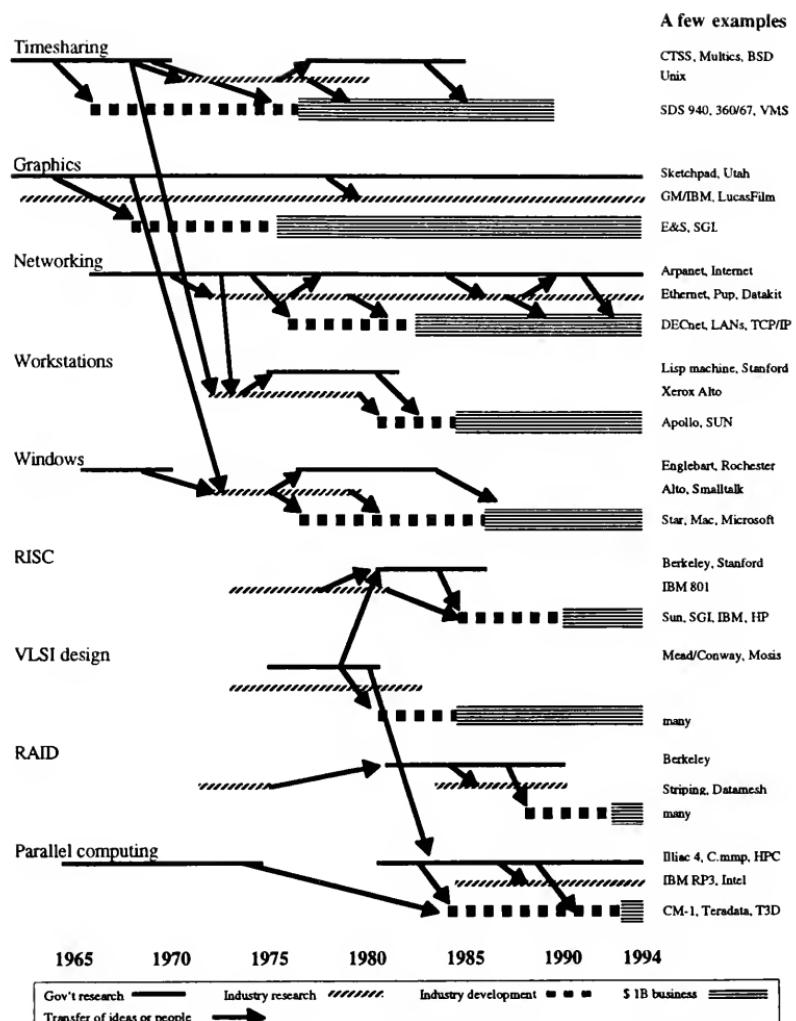


FIGURE 1.2 Government-sponsored computing research and development stimulates creation of innovative ideas and industries. Dates apply to horizontal bars, but not to arrows showing transfer of people. Table 1.1 is a companion to this figure.

**BOX 1.3 An Example of the Impact of Federal
R&D Support in Establishing a Field: Computer Graphics**

In the middle 1960s, using a computer on loan from the Air Force and financial support from the Central Intelligence Agency, the computer graphics group at Harvard University demonstrated a prototype virtual reality system. This work contributed significantly to the technological and personnel foundation for the Evans and Sutherland Corporation, which subsequently provided computer-based equipment for training pilots—equipment that is used today by the U.S. military and by commercial pilots the world over.

In the early 1970s the University of Utah was host to a leading program in computer graphics. Dave Evans went there to found a computer science department, and realizing that his department could not be all things to all people, he specialized in computer graphics. ARPA provided the main research support.

At that time nearly all pictures of three-dimensional objects were drawn with lines only. The resulting images appeared to be of wire frames. They were not very realistic. The Utah group worked mainly on techniques for increasing the realism by omitting parts of objects that were hidden behind other parts and by shading the surfaces of the objects. The resulting pictures were much more realistic.

Two key developments had particularly significant impact. First, Watkins and others, following a suggestion of Evans, developed a set of incremental techniques for computing what parts of an object were hidden. The key observation was that two parts of the image must be nearly the same if they are close together. When some part of an image has been computed, nearby parts are easier to compute than they would be if computed in isolation.

Second, Gouraud and Phong and others developed incremental algorithms for shading solid surfaces. Prior to their work the best images appeared to be made with flat surfaces; each surface was painted a single shade according to the angle between it, the light, and the observer. Many workers sought methods for representing objects with curved surfaces, but it was then and is still difficult. Instead, Gouraud invented a trick. He painted each surface a gradually changing shade in such a way that at the joints between surfaces they had the same shade. With no demarcation line, the human eye thinks the resulting surface is smooth even though it is made of little flat plates. Phong went a step further, computing highlights as if the surface were curved. Gouraud shading and Phong shading are in standard use everywhere today. It is particularly interesting to note that when government support started, no one knew that these technologies were possible and the people who made the key discoveries were not yet involved.

The vast implications of computer graphics (what-you-see-is-what-you-get document creation systems, scientific visualization, the entertainment industry, virtual reality) were of course totally unforeseen at the time that this fundamental research was undertaken. In addition to the specific developments cited above, an essential contribution was the many individuals whose training in universities benefited from ARPA support. A few of the more prominent are John Warnock of Adobe Systems (\$300 million per year), Jim Clark (formerly) of Silicon Graphics Inc. (\$1 billion per year), Henry Fuchs of the University of North Carolina, and Ed Catmull of Pixar. Many others carried the knowledge to companies and academic institutions throughout the nation.

- *Even for defense applications, supporting research on strategically motivated but broadly applicable computing and communications technology has clearly been the right approach.* In the past, many defense applications and requirements presaged commercial applications and requirements. Today, commercial computer systems and applications often find use in defense environments.
- *Research and development take time.* At least 10 years, more often 15, elapse between initial research and commercial success. This is true even for research of strategic importance. And it is true in spite of the rapid pace of today's product development, as indicated in the timeline for recent commercial successes such as RISC and windows (see Figure 1.2).

BOX 1.4 Federal R&D Support Propels Database Technology

The database industry generated about \$7 billion in revenue in 1994 and is growing at a rate of 35 percent per year. All database industry leaders are U.S.-based corporations: IBM, Oracle, Sybase, Informix, Computer Associates, and Microsoft. In addition, there are two large specialty vendors, both also U.S.-based: Tandem, selling over \$1 billion per year of fault-tolerant transaction processing systems, and AT&T-Teradata, selling about \$500 million per year of data mining systems. In addition to these well-established companies, there is a vibrant group of small companies specializing in application-specific databases, object-oriented databases, and desktop databases.

A very modest federal research investment, complemented by a modest industrial research investment, led directly to U.S. dominance of this market. It is not possible to list all the contributions here, but three representative research projects are highlighted that had major impact on the industry.

1. Project Ingres started at the University of California, Berkeley, in 1972. Inspired by Codd's landmark paper on relational databases, several faculty members (Stonebraker, Rowe, Wong, and others) started a project to design and build a system. Incidental to this work, they invented a query language (QUEL), relational optimization techniques, a language binding technique, and interesting storage strategies. They also pioneered work on distributed databases.

The Ingres academic system formed the basis for the Ingres product now owned by Computer Associates. Students trained on Ingres went on to start or staff all the major database companies (AT&T, Britton Lee, HP, Informix, IBM, Oracle, Tandem, Sybase). The Ingres project went on to investigate distributed databases, database inference, active databases, and extensible databases. It was rechristened Postgres, which is now the basis of the digital library and scientific database efforts within the University of California system. Recently, Postgres spun off to become the basis for a new object-relational system from the start-up Illustra Information Technologies.

2. System R was IBM Research's response to Codd's ideas. His relational model was at first very controversial; people thought that the model was too simplistic and that it would never perform well. IBM Research management took a gamble and chartered a small (10-person) effort to prototype a relational system based on Codd's ideas. That effort produced a prototype, System R, that eventually grew into the DB2 product series. Along the way, the IBM team pioneered ideas in query optimization, data independence (views), transactions (logging and locking), and security (the grant-revoke model). In addition, the SQL query language from System R was the basis for the ANSI/ISO standard.

The System R group went on to investigate distributed databases (project R*) and object-oriented extensible databases (project Starburst). These research projects have pioneered new ideas and algorithms. The results appear in IBM's database products and those of other vendors.

3. The University of Wisconsin's Gamma system was a highly successful effort that prototyped a high-performance parallel database system built of off-the-shelf system components.

During the 1970s there had been great enthusiasm for database machines—special-purpose computers that would be much faster than general-purpose systems running conventional databases. These research projects were often based on exotic hardware like bubble memories, head-per-track disks, or associative random access memory. The problem was that general-purpose systems were improving at a rate of 50 percent per year, and so it was difficult for exotic systems to compete with them. By 1980, most researchers realized the futility of special-purpose approaches, and the database machine community switched to research on using arrays of general-purpose processors and disks to process data in parallel.

The University of Wisconsin was home to the major proponents of this idea in the United States. Funded by government and industry, researchers prototyped and built a parallel database machine called Gamma, whose hardware base was an early Intel scalable parallel machine. That system produced ideas and a generation of students who went on to staff all the database vendors. Today, the highly successful parallel database systems from IBM, Tandem, Oracle, Informix, Sybase, and AT&T all have a direct lineage from the Wisconsin research on parallel database systems. The use of parallel databases systems for data mining is now the fastest-growing component of the database server industry.

The Gamma project evolved into the Exodus project at Wisconsin (focusing on an extensible object-oriented database). Exodus has now evolved to the Paradise system, which combines object-oriented and parallel database techniques to represent, store, and quickly process huge Earth-observing satellite databases.

SOURCE: James Gray and others for the Computing Research Association; reproduced with permission.

CONTINUED FEDERAL INVESTMENT IS NECESSARY TO SUSTAIN OUR LEAD

What must be done to sustain the innovation and growth needed for enhancing the information infrastructure and maintaining U.S. leadership in information technology? Rapid and continuing change in the technology, a 10- to 15-year cycle from idea to commercial success, and successive waves of new companies are characteristics of the information industry that point to the need for a stable source of expertise and some room for a long-term approach. Three observations seem pertinent.

1. *Industrial R&D cannot replace government investment in basic research.* Very few companies are able to invest for a payoff that is 10 years away. Moreover, many advances are broad in their applicability and complex enough to take several engineering iterations to get right, and so the key insights become "public" and a single company cannot recoup the research investment. Public investment in research that creates a reservoir of new ideas and trained people is repaid many times over by jobs and taxes in the information industry, more innovation and productivity in other industries, and improvements in the daily lives of citizens. This investment is essential to maintain U.S. international competitiveness.

Of course, industrial R&D also contributes to the nation's pool of new ideas, but a company may postpone exploiting its ideas if they disturb existing business. A good example is the evolution of RISC processors shown in Figure 1.2. RISC was invented by John Cocke, an IBM researcher, but IBM made no RISC products for a decade. Only after the ideas were embraced and extended in government-sponsored work at universities did industry adopt them, and this adoption was initiated by young companies, including Sun Microsystems, and start-ups, including MIPS. Now, a decade later, IBM is one of the leaders in exploiting RISC technology, but the cost to IBM of this delay has been significant. Firms have regularly failed to adapt to change as evidenced by the departure from the computer business of GE, RCA, Honeywell, Philco, Perkin-Elmer, Control Data, and Prime; the folding together by merger of other manufacturers; and major downsizing at IBM and DEC. It often is easier for a start-up to form, raise venture capital, and succeed than for an established firm to abandon a currently successful direction in favor of a new approach just when the old approach is most financially successful. Even in a vigorous industrial R&D climate, then, federal investment in research is necessary, both for its long-term focus and for its ability to incubate ideas to the point of clear commercial viability.

But the need for federal investment in research is further compounded by the fact that industrial R&D is already under stress. In the computing hardware and software sector, for example, although a small number of new R&D enterprises have been launched, most notably by Microsoft, there has been a general consolidation of resources by companies such as IBM and DEC, including an apparent reduction in their research effort or at least a greater emphasis on short-term R&D—a change in emphasis is evident to insiders and close observers but not easy to document.⁵ The industry-wide level of R&D as a percentage of sales has also been brought down by the tendency of low-price vendors, such as Dell and Gateway, to ride on the research conducted by others.

The trend toward reduced industrial R&D appears also in the telecommunications industry. The 1984 divestiture of AT&T led to a smaller Bell Laboratories and to the creation of Bell Communications Research (Bellcore), a shared research facility for the seven regional Bell holding companies. Recent deregulation has encouraged a reduction of basic research at both AT&T and Bellcore. Lacking significant research capability at its individual service companies, the cable television industry depends on research done by its

hardware vendors and its shared CableLabs. Although more new technology has been deployed in telecommunications since deregulation in the early 1980s, and although in both computing and communications there are more companies selling products now than there were 15 years ago, today's sales are based on yesterday's research and do not guarantee a sufficient foundation for tomorrow's sales. Competition in an industry can promote technological growth, but competition alone is not the source of innovation and leadership.

Because of the long time scales involved in research, the full effect of decreasing investment in research may not be evident for a decade, but by then, it may be too late to reverse an erosion of research capability. Thus, even though many private-sector organizations that have weighed in on one or more policy areas relating to the enhancement of information infrastructure typically argue for a minimal government role in commercialization, they tend to support a continuing federal presence in relevant basic research.⁶

2. *It is hard to predict which new ideas and approaches will succeed.* Over the years, federal support of computing and communications research in universities has helped make possible an environment for exploration and experimentation, leading to a broad range of diverse ideas from which the marketplace ultimately has selected winners and losers. As demonstrated by the unforeseen applications and results listed in Table 1.1, it is difficult to know in advance the outcome or final value of a particular line of inquiry. But the history of development in computing and communications suggests that innovation arises from a diversity of ideas and some freedom to take a long-range view. It is notoriously difficult to place a specific value on the generation of knowledge and experience, but such benefits are much broader than sales of specific systems.

3. *Research and development in information technology can make good use of equipment that is 10 years in advance of current "commodity" practice.* When it is first used for research, such a piece of equipment is often a supercomputer. By the time that research makes its way to commercial use, computers of equal power are no longer expensive or rare. For example, the computer graphics techniques that are available on desktop workstations today, and that will soon be on personal computers and set-top boxes, were pioneered on the multimillion-dollar supercomputers of the 1960s and 1970s. Part of the task in information technology R&D is to find out how today's supercomputers can be used when they are cheap and widely available, and thus to feed the industries of tomorrow.

The large-scale systems problems presented both by massive parallelism and by massive information infrastructure are additional distinguishing characteristics of information systems R&D, because they imply a need for scale in the research effort itself. In principle, collaborative efforts might help to overcome the problem of attaining critical mass and scale, yet history suggests that there are relatively few collaborations in basic research within any industry, and purely industrial (and increasingly industry-university or industry-government) collaborations tend to disseminate results more slowly than university-based research.

The government-supported research program (on the order of \$1 billion for information technology R&D) is small compared to industrial R&D (on the order of \$20 billion; Coy, 1993), but it constitutes a significant portion of the research component, and it is a critical factor because it supports the exploratory work that is difficult for industry to afford, allows the pursuit of ideas that may lead to success in unexpected ways, and nourishes the industry of the future, creating jobs and benefits for ourselves and our children. The industrial R&D investment, though larger in dollars, is different in nature: it focuses on the near term—increasingly so, as noted earlier—and is thus vulnerable to major opportunity costs. The increasing tendency to focus on the near term is

affecting the body of the nation's overall R&D. Despite economic studies showing that the United States leads the world in reaping benefits from basic research, pressures in all sectors appear to be promoting a shift in universities toward near-term efforts, resulting in a decline in basic research even as a share of university research (Cohen and Noll, 1994). Thus, a general reduction in support for basic research appears to be taking place.

It is critical to understand that there are dramatic new opportunities that still can be developed by fundamental research in information technology—opportunities on which the nation must capitalize. These include high-performance systems and applications for science and engineering; high-confidence systems for applications such as health care, law enforcement, and finance; building blocks for global-scale information utilities (e.g., electronic payment); interactive environments for applications ranging from telemedicine to entertainment; improved user interfaces to allow the creation and use of ever more sophisticated applications by ever broader cross sections of the population; and the creation of the human capital on which the next generation's information industries will be based. Fundamental research in computing and communications is the key to unlocking the potential of these new applications.

How much federal research support is proper for the foreseeable future and to what aspects of information technology should it be devoted?⁷ Answering this question is part of a larger process of considering how to reorient overall federal spending on R&D from a context dominated by national security to one driven more by other economic and social goals. It is harder to achieve the kind of consensus needed to sustain federal research programs associated with these goals than it was under the national security aegis. Nevertheless, the fundamental rationale for federal programs remains (Cohen and Noll, 1994, p. 73):

That R&D can enhance the nation's economic welfare is not, by itself, sufficient reason to justify a prominent role for the federal government in financing it.

Economists have developed a further rationale for government subsidies. Their consensus is that most of the benefits of innovation accrue not to innovators but to consumers through products that are better or less expensive, or both. Because the benefits of technological progress are broadly shared, innovators lack the financial incentive to improve technologies as much as is socially desirable. Therefore, the government can improve the performance of the economy by adopting policies that facilitate and increase investments in research.

TODAY THE HPCCI IS THE UMBRELLA FOR MOST GOVERNMENT-SPONSORED COMPUTING AND COMMUNICATIONS RESEARCH

Today, the HPCCI is the umbrella sheltering most government-sponsored computing and communications research. The HPCCI is thus responsible for sustaining the nation's lead in information technology. It centers on rising performance as the driver for progress across a wide range of technologies.

"High performance" means bringing a large quantity of computing and communications to bear on one problem. It is far broader than "supercomputing," which was the focus of early public policy in this area. It is also a moving target—the threshold for what is considered "high performance" advances, as ever-increasing performance levels become more broadly available. The supercomputer of this generation is the group server of the next generation and the personal computer of the generation after that. The same is true for communications: today's leading edge is tomorrow's mainstream.

Focusing research on the leading edge of performance accelerates the broad availability of what starts out as limited-access technology, in the following ways:

- *By advancing the hardware and software systems themselves.* Many techniques now used to build mainstream computers and their software were originally developed for high-performance computing: specialized floating-point processing, pipelining, multiprocessing, and multiple instruction issue, among others;
- *By creating new applications today so that they will be mature when the hardware that can run them is cheap and ubiquitous.* It can take much longer to develop and fully exploit a new application than to build a new computer. Overcoming this lag is one of the drivers of work on the Grand and National Challenge concepts; and
- *By attacking problems that would otherwise be beyond reach for several years, thus speeding up the development of new fields of science, engineering, medicine, and business.* National access to machines with 100 to 1,000 times the memory and speed of researchers' desktop machines allows them to make qualitative jumps to exploring frontier research problems of higher dimensionality, greater resolution, or more complexity than would otherwise be possible.

Fundamental but strategic research under the HPCCI—which now encompasses most of the academic computing research sponsored by the federal government—creates a strong pull on the computer science and engineering research community, the user community, and the hardware, software, and telecommunications vendors. For example, it was evident to individuals in the computing and communications research community that as VLSI circuit technology developed, it would favor computing structures based on the large-scale replication of modest processors, as opposed to the small-scale replication of processors of the highest attainable individual performance. (The highest-speed circuits are expensive to design, produce, maintain, and operate.) This vision of high-performance computing and communications based on parallelism brought three major technical challenges: (1) interconnection and memory architecture—how to unite large numbers of relatively inexpensive processors into systems capable of delivering truly high performance, and (2) programming—how to program such collections of processors to solve large and complex problems. In the networking arena, the obvious issues of large-scale, widespread use and high-speed transmission were compounded by added problems of connecting heterogeneous systems and achieving high reliability.

The technical and economic imperatives that led to the HPCCI are discussed in some detail in Appendix A. HPCCI was, and continues to be, an appropriate thrust. As Chapter 2 discusses, the HPCCI can boast a broad range of very significant accomplishments. However, the committee sees an unhealthy dependence of our nation's critical leadership in information technology on the fate of a single initiative. First, not all important computing research is focused on high performance, although the politics of funding have encouraged researchers and agencies to paint everything with an HPCCI brush. Second, we cannot afford to cripple computing research if the nation's attention and resources turn away from any single goal. At the beginning of the HPCCI in 1992, its increasing momentum made association with it attractive and helped the initiative attain both intellectual and financial critical mass. The rise of the National Information Infrastructure initiative, though, underscores how changeable the federal funding process may be.

We must move toward a more mature approach in which a substantial focus on goals of obvious national importance is combined with a diversified program of long-term exploration of important research problems that support the strategic information technology industry. We can change the HPCCI's name, we can change its orientation, but we must move forward. Continuing the momentum of this successful initiative is essential to ongoing U.S. prosperity and leadership in information technology.

NOTES

1. Microcomputers (personal computers) are defined as computers with a list price of \$1,000 to \$14,999; see CBEMA (1994), pp. 60-61. Forrester (1994, pp. 2-3) estimates the share of households with PCs at about 20 percent, based on their survey of households and Bureau of Census data. Forrester's model accounts for retirements of older PCs and for households with multiple PCs. This is a lower estimate than the Software Publishing Association's widely cited 30 percent share. Building on an unusual fourth quarter sales surge, almost 7 million PCs were sold for residential use in 1994 (Markoff, 1995). According to another source, fourth quarter 1994 U.S. PC shipments were 32 percent higher than the corresponding quarter in 1993, 5.8 million units, feeding a 27 percent surge in worldwide PC shipments for 1994 in total (Carlton, 1994).

2. According to Roach (1994, p. 12), "IT [information technology] expenditures now comprise fully 45 percent of total business outlays on capital equipment—easily the largest line item in Corporate America's investment budget and up dramatically from a 20% share seen as recently as 1980. . . . But does the technology story have staying power? We believe that the new dynamics of technology demand should continue to power the U.S. economy throughout the 1990s. Indeed, *the technology-capital-spending link is an integral element of the productivity-led recovery scenario that lies at the heart of our basic macro call for the United States* [emphasis in original]. In this light, IT is the principal means by which businesses can improve upon the efficiency gains first derived from cost-cutting, facilitating the transition between slash-and-burn downsizing and the rebuilding eventually required for sustained competitive prowess. Without increasing emphasis on technology, the economy gets hollowed out. The good news is that the long-awaited technology payback suggests this darker scenario won't come to pass."

3. See U.S. DOC (1994); the Department of Commerce utilizes data from the U.S. Bureau of the Census series, the *Annual Survey of Manufactures*. It places the value of shipments for the information technology industry at \$421 billion for 1993. This number omits revenue from equipment rentals, fees for after-sale service, and markups in the product distribution channel. It also excludes office equipment in total. It includes computers, storage devices, terminals, and peripherals; packaged software; computer program manufacturing, data processing, information services, facilities management, and other services; and telecommunications equipment and services.

See also CBEMA (1994); CBEMA values the worldwide 1993 revenue of the U.S. information technology industry at \$602 billion. In addition to including office equipment, it shows larger revenues for information technology hardware and telecommunications equipment than does the Department of Commerce.

4. In addition, European and Japanese manufacturers have significant sales in computer-related products such as telecommunications switches, semiconductors, and even high-performance computing equipment.

5. See Rensberger (1994) and Corcoran (1994). Also Coy (1993) indicates that data on individual information technology companies show 1992 R&D spending as a percent of sales ranging from 1 to 15 percent, concentrated at the lower end of that range; industry segment statistics fall in the same range. The figures suggest that despite the "high-tech" image of the industry, less R&D is conducted than many believe, and many firms capitalize on research conducted by others.

6. See, for example, CSPP (1994), pp. 1-2. A broad argument for a federal role in support of basic research in critical technologies, including computing and communications, is presented in a Council on Competitiveness (1991) report.

7. A point of reference can be found in a mid-1994 document from the Electronics Subcommittee of the National Science and Technology Council (NSTC, 1994a). It calls attention to U.S. dependence on a healthy electronics industry and speaks of efforts to work with industry to "develop a roadmap for electronics that will illuminate gaps in government-sponsored research and infrastructure efforts," focusing on "information products that connect to information networks, including the National Information Infrastructure (NII)"

The High Performance Computing and Communications Initiative

The High Performance Computing and Communications Initiative (HPCCI) has been the focal point of federal support for U.S. computing and communications research and development since 1989. It became official in 1991 with Office of Science and Technology Policy (OSTP) support and enactment of the High Performance Computing Act of 1991. It includes five programs: Advanced Software Technology and Algorithms, Basic Research and Human Resources (although there is basic research in the other four programs also), High-Performance Computing Systems, National Research and Education Network (NREN), and since FY 1994, Information Infrastructure Technology and Applications (IITA).¹ Appendix A outlines the origins and early history of the HPCCI, including an explanation of associated technology trends and indications of evolution of the initiative's emphases. This chapter discusses the HPCCI's goals and contributions to date and identifies key substantive and practical issues to be considered as the initiative evolves.²

HPCCI: GOALS AND EMPHASES

The HPCCI has several broad goals (NCO, 1993):

- Extend U.S. leadership in high-performance computing and networking technologies;
- Disseminate the technologies to accelerate innovation and serve the economy, national security, education, and the environment; and
- Spur gains in U.S. productivity and industrial competitiveness.

Because these goals relate advances in computing and communications technologies to the achievement of benefits from their use, the HPCCI has from its inception provided for the joint advancement of technologies and applications. The HPCCI has pursued several specific strategic objectives.

Basic Objectives

Teraflop Capability

The specific objective for computer development was to develop teraflop capability by the mid-1990s.³ This objective was comparable to two that had been achieved earlier by forerunners of today's computer science community at the request of the federal government: peak available

computing power was increased by several orders of magnitude during World War II, when federal interests in cryptanalysis and other wartime needs led to the development of the vacuum tube computer, and in the late 1950s, when federal interests in military command and control led to transistorized computers.⁴ As economist Kenneth Flamm has observed, "By tracking the origins and history of key pieces of technology, a simple but important point can be established: at certain crucial moments in history, private commercial interests, in the absence of government support, would not have become as intensely involved in certain long-term basic research and radical new concepts" (Flamm, 1988, p. 3). The teraflop objective has inspired parallel multi-microprocessor computers as the means for providing the next major jump in computer power.⁵

The teraflop objective has generated both attention and misunderstanding. Progress required building a number of machines large and fast enough to reward software researchers and application users with major gains in performance, thereby motivating them to develop the code that could make the high-performance machines useful. (See Appendix A for more information on the development of high-performance hardware and software and their interaction.) The costliness of this undertaking, compounded by the highly publicized financial difficulties of two entrepreneurial ventures, Thinking Machines Corporation (TMC) and Kendall Square Research (KSR), aimed at commercializing massively parallel computing systems, attracted criticism of the HPCCI.

However, that criticism appears largely misdirected. First, entrepreneurial ventures are always risky, and the two in question suffered from managerial weakness at least as much as questionable technology choices.⁶ Contemporaneously, more established firms (e.g., Cray Research, IBM, Intel Supercomputing, Convex Computer, and Silicon Graphics Inc.; Parker-Smith, 1994a) have persevered, and others (e.g., Hitachi and NEC in Japan; Kahaner, 1994b, and Parker-Smith, 1994b) have entered or expanded their presence in the parallel systems market. Second, focusing attention on the high initial costs for stimulating development and use of parallel processing systems detracts from the achievement of successful proofs of concept and dissemination of new approaches to computation.

Although the teraflop objective was ambitious for the time scale set, it was intended as a driver and thus is best viewed as indicating a direction, not a destination; the need for progress in computing will continue beyond the teraflop capability.⁷ In that respect, its appropriateness was affirmed by the 1993 Branscomb panel.⁸ The teraflop objective has, in fact, served to focus attention on the task of combining and harnessing vast amounts of computer power from many smaller computers. The technology is now sufficiently developed that a teraflop machine could be realized today, although exactly when to do so should be left to the economics of users and their applications.⁹

High-speed Networks

Another direction-setting objective of the HPCCI was the achievement of data communications networks attaining speeds of at least 1 gigabit per second. Although by the mid-1980s major telecommunications networks already had gigabit-plus trunk circuits in their backbones, the HPCCI was intended to lead to much broader deployment of and access to gigabit-speed networks connecting general-purpose computers.¹⁰ This objective drove progress in switching, computing hardware and software, interfaces, and communication protocols.¹¹ (See Appendix B.)

Grand Challenges

A third original objective related to applications of high-performance computing and communications technologies: to define and attack Grand Challenge problems. High-performance

computing and communications national centers in the various agencies already were providing access to tens of thousands of researchers in hundreds of institutions when the HPCCI began. Drawing from this national base of users, various HPCCI agencies have defined a series of Grand Challenge problems (see Appendix D for list) and chosen teams to attack them. The Grand Challenge teams are typically both interdisciplinary and multi-institutional. The scientific problems are picked for their intrinsic scientific merit, the need for high-performance computation and communications, the opportunities for synergistic interaction of computer scientists with computational scientists, and the scientific and societal benefits to be gained from their solution. For example, better weather prediction involves solving massive sets of equations, experimenting with models, and comparing the results obtained with them to increasingly large volumes of data collected by weather-monitoring instruments.¹² High-performance computing provides the faster model computation essential to timely assessment of a sufficiently large volume of alternative weather patterns for a given period (e.g., a month).¹³ The results include not only greater scientific understanding, but also the benefits to businesses, individuals, and governments that come from faster, more accurate, and more detailed forecasts.

Expanded Objectives

The set of HPCCI objectives has been expanded through legislative and agency activities. The High Performance Computing Act of 1991, Public Law 102-194, broadened the applications concerns to include the so-called National Challenges—explorations of high-performance computing and communications technology for applications in such areas as education, libraries, manufacturing, and health care. PL 102-194 also reinforced the communications aspects of the HPCCI, elaborating the concept and objectives for the NREN program and emphasizing networking applications in education. Officials involved with the HPCCI have noted that although PL 102-194 was never complemented by specific appropriations legislation, its principles have driven HPCCI activities in relevant agencies, including early explorations relating to National Challenges and the formation of the Information Infrastructure Technology and Applications (IITA) component in FY 1994.¹⁴ The National Challenges, IITA, and the network aspects of PL 102-194 also included attention to short-term and practical concerns (e.g., expanding access to technology facilities and capabilities), complementing the long-term, basic research problems that remain at the heart of HPCCI.¹⁵

The tension between long term and short term, between basic research and applications, is fundamental to the differences in opinion voiced about the HPCCI and its merits, accomplishments, and desired directions. Based on its direct observations of work funded under the HPCCI and on its discussions with others in universities, industry, and the government, the committee affirms the value of the basic research associated with the HPCCI, research that is informed by needs associated with important applications of national interest.

A fundamental issue shaping the evolution of the HPCCI is the balance to be struck between the support of applications that use high-performance computing and communications technologies and the support for computer science research on new high-performance computing and communications technologies.¹⁶ The committee's analysis of the FY 1995 HPCCI budget request (Appendix C) shows that out of the total request of \$1.15 billion, \$352 million (30 percent) would be invested in basic research in computer, software, or communication technologies; \$205 million (18 percent) in applied computer science research in common applications support, artificial intelligence, and human-machine interface; \$176 million (15 percent) in direct support of applications and computational science; and \$297 million (26 percent) in supercomputing and communications infrastructure. It is hard to interpret these statistics, however, without an understanding of the nature and the value of the work labeled "applications." The HPCCI has been aimed at catalyzing a paradigm shift, which involves the synergistic interaction of people developing

the technology and people using the technology.¹⁷ The HPCCI includes mission-related activities that may drive computing and communications research and development (R&D) and/or applications that call for significant technology development.

Within the computer science and engineering field, there has been considerable debate over the degree to which computing research should be driven by applications concerns as opposed to intrinsic computer science concerns, given that both approaches to research have historically yielded considerable spinoffs to other sciences and the economy.¹⁸ To computer scientists and engineers, HPCCI is viewed as the first major federal initiative that emphasizes the science of computing and communications, which is addressed in conjunction with exploration of problems involving other fields of science and engineering, loosely aggregated as computational science. To computational scientists, the emphasis is predictably on the problems in their domains and on the difficulty of developing appropriate domain-specific computational techniques. These differences in outlook result in differences in what each community calls an "application," as well as differences in requirements for R&D.

HPCCI ACCOMPLISHMENTS

Accomplishments under the HPCCI to date reveal two key trends: better computing and computational infrastructure and increasing researcher-developer-user synergy. In the committee's expert judgment, HPCCI has been generally successful. That assessment is necessarily qualitative and experiential, because it is too early yet to observe the full impact of the initiative.

The Issue of Measurement

Early measurement of the impact of HPCCI research is problematic. As Chapter 1 points out, the time for progress from idea to product involves a decade or more, well beyond a single fiscal year. Independent of impact, individual projects may take a few years simply to reach completion.¹⁹ Consequently, the accomplishments of the HPCCI are only just becoming apparent.

Moreover, it is difficult to evaluate early on how good individual ideas are and what their worth may prove to be. Many researchers have expressed concern that the push for immediately measurable results has led to an unrealistic emphasis on short-term gains and has diverted efforts from conducting productive research to maintaining "paper trails."²⁰ However, the pressures on agencies to maximize the return on limited research funds seems to discourage funding of more innovative—and therefore riskier—exploration that may not necessarily succeed (Rensberger, 1994). The problem of measurement is compounded by the fact that a considerable amount of HPCCI research addresses enabling technologies whose benefits or outcomes may be evident only indirectly.

How best to assess results is unclear—key questions include the kinds of reviews already undertaken by agencies and with what effect; how evaluations based on outside expertise should be combined with in-house agency know-how; whether to focus on reviewing progress for a program as a whole or progress in individual grants; the costs in time and money of different approaches and comparison of the benefits in terms of review quality, scope, and timeliness to the costs; and so on. The committee recognizes that data and analysis are needed to support decision making about any new approaches to evaluation; it did not have the time or resources to pursue such analysis.²¹ Moreover, complementing the committee's observation that much of the evidence on outcomes is anecdotal is a recent National Research Council study pointing out that good, relevant data (on scientific research in general) are hard to find and even harder to draw inferences from (CPSMA, 1994).

Better Computing and Computational Infrastructure

The HPCCI has contributed substantially to the development, deployment, and understanding of computing and communications facilities and capabilities as infrastructure. It has helped transform understanding of how to share resources and information, generating proofs of concept and understanding that are of value not only to the scientific research community but also to the economy and society at large.

The HPCCI has directly strengthened academic computing and computational infrastructure, building on the National Science Foundation's (NSF's) significant investments in university computing infrastructure over more than a decade.²² The NSF infrastructure program has stimulated complementary investments by other federal agencies, industry, and universities themselves—an impact that, like other HPCCI contributions to stimulating a growing foundation of activity, is difficult to assess directly. This academic base, in particular, academic research in experimental computer science and engineering,²³ is fundamental to the development and application of high-performance and other computing and communications technologies (CSTB, 1994a).

By providing access (often over the Internet) to state-of-the-art computer resources and to expertise to help researchers learn how to use them, the HPCCI has also enabled research in a wide range of science and engineering disciplines to be performed that would not otherwise have been possible. Appendix D lists relevant examples from the Grand Challenge activities, and Appendix E points out instances related to the NSF supercomputer center activities, which fall under the HPCCI umbrella despite having some separate roots.

Within the NREN program, NSFNET and other components such as ESNet and the NASA Science Internet have helped to extend networking across the science research community (CSTB, 1994d). Through the internetworking provided by the Internet, connectivity and experimentation with network-based infrastructure have begun to spread rapidly beyond the research community into primary and secondary education, government, industry, and other elements of the private sector (CSTB, 1994d). The Internet has demonstrated the value of widespread access to a common, sophisticated, and increasingly integrated technology base, and it illustrates how a small investment by the federal government can be highly leveraged by additional investments from industry.²⁴

The HPCCI approach to developing high-performance computing and communications infrastructure has been affirmed in similar steps taken recently by the Japanese government and industry. David Kahaner of the Office of Naval Research has chronicled Fujitsu's progress in developing parallel processing technology, noting its establishment of research facilities providing worldwide access to its systems in order to obtain the large user base needed to refine its hardware designs and, in particular, to develop the software and applications required to make systems successful (Kahaner, 1994a). Kahaner has also reported on Japanese plans and progress for upgrading high-performance capabilities in public institutions, noting, among other things, the Japanese government's increasing emphasis on basic research.²⁵

Increasing Researcher-Developer-User Synergy

The HPCCI has fostered productive interactions among the researchers and developers involved in creating high-performance computing and communications technology and those who use this technology in their own work, most notably computational scientists, but also a broad spectrum of other users. Building on the varying needs and perspectives of the three groups, complex problems are being solved in unique ways.

In particular, the HPCCI has funded cross-disciplinary teams associated with the Grand Challenge projects to solve complex computational problems and produce new software for the new

parallel systems. These teams interact with hardware and operating system/compiler researchers and developers to address complex problems through use of the largest computer systems, including those housed at the NSF supercomputer centers. Their work has provided vendors with key insights into the limitations of their architectures.²⁶ Although these users' requirements are more specialized than those typical of the commercial market for parallel systems, such collaborative work has contributed to enhancing the development and application of high-performance computing and communications technologies. For example, astrophysicists' work on problems in cosmology has stimulated improved handling of fast Fourier transforms in high-performance system compilers that has also benefited commercial applications of seismology in oil exploration.²⁷

Like collaboration in other areas, that between computer and computational scientists has not always come easily. In particular, there has been some controversy concerning the relative emphasis on advancing disciplinary knowledge, on the one hand, and advancing the state of the art in high-performance computing and communications, on the other. Nevertheless, the HPCCI has provided a structure and a set of incentives to foster collaborations that many computational scientists believe would not be supported under programs aimed at nurturing individual disciplines.²⁸

Impact of Broad Collaboration

Many notable HPCCI accomplishments are the result of broad collaborations. In many instances, they build on foundations that predated the HPCCI, although HPCCI funding, facilities, and focus may have provided the push needed for their realization. The Mosaic browser (Box 2.1) epitomizes both the cumulative nature and broad impact of the development of technologies associated with the HPCCI.

- The HPCCI has driven progress on Grand Challenge problems in disciplines such as cosmology, molecular biology, chemistry, and materials science. Parallel computation has enhanced the ability to do rapid simulations in science and engineering (Kaufmann and Smarr, 1993). Recognition of this development continues to spread across the research community.
- The HPCCI has furthered the development of new modes of analyzing and/or visualizing complex data and in many cases has contributed to more effective interworking between supercomputers and desktop graphics workstations. Visualizations of the numerical output of the largest computers require specialized graphics computers, whose speed would have made them supercomputers in their own right a few years ago. Examples include visualization of complex motions of large biomolecules, intricate engineering assemblies, and the large-scale structure in the universe.
- The HPCCI has made parallel computing widely accepted as the practical route to achieving high-performance computing, as can be seen in the recent growth in sales of parallel systems.²⁹ Although the market for larger-scale parallel-processing systems is inherently small, it is nevertheless growing. Box 2.2 gives a few of many possible examples of the applications being developed.

BOX 2.1 Mosaic and the World Wide Web

The development in 1993 of the National Center for Supercomputing Applications (NCSA) Mosaic browser shows how the HPCCI has been able to create successful new applications enabled both by new capabilities and by prior developments in information technology.

The forerunner of the Internet (ARPANET) was developed in the late 1960s to link computers and scientists performing defense-related research throughout the United States. By the time of the HPCCI's formal initiation in FY 1992, the Internet had become the most popular network linking researchers and educators at the post-secondary level throughout the world. The development of gopher at the University of Minnesota in the early 1990s was a key step in establishing the Internet as an information resource that could be used through a consistent user interface. At about the same time, researchers at the European Laboratory for Particle Physics, CERN, had developed and implemented the World Wide Web (WWW), a network-based hypertext system that allowed the linking of one piece of information to another across the network. Users accessed WWW information through "browsers" that allowed them to activate a hypertext link in a document to retrieve and display related information regardless of its physical location. Early browsers were text-based, presenting the user with a menu of numbered choices, whereas slightly later browsers made use of the "point-and-click" capabilities of the mouse within a graphical user interface. The WWW and its browsers sought to present users with a consistent interface with which to access all existing methods of Internet navigation and information retrieval.

Meanwhile, the HPCCI had provided funding for research into advanced networking technologies and for the deployment of a high-capacity backbone, enabling the rapid transfer of large amounts of data across the network. In 1993, software developers at the NCSA, one of the centers supported by HPCCI funds from the National Science Foundation, developed an easy-to-use graphical browser for the WWW known as NCSA Mosaic, or sometimes simply Mosaic. It allowed the inclusion of images in WWW documents and even allowed the images themselves to be links to other information. Continuing development of Mosaic enabled the inclusion of audio and video "clips" within hypermedia documents. By November 1993, Mosaic browsers were available for the three most popular computer operating environments: Apple Macintosh, Microsoft Windows, and X Window on UNIX platforms. One year later, users have downloaded more than 1 million copies of Mosaic software, NCSA's scalable WWW server (the world's busiest) is handling over 4 million connections per week, and Mosaic is credited by many for the current and dramatic surge in use of and interest in the Internet.

Perhaps even more significantly, Mosaic has served as the genesis of a wide range of commercial developments. The University of Illinois, which owns the intellectual property associated with Mosaic, has named Spyglass, Inc. as the master sublicensee for the software. So far over 20 companies have licensed Mosaic, creating over 12 million commercially licensed copies. In addition, other companies such as Netscape, IBM, Pipeline, Booklink, MCC, and NetManage have created alternative WWW browsers.

Many other entities have become information providers (a 100-fold increase in WWW servers has occurred in the last 2 years), and new security additions to the underlying Mosaic/WWW infrastructure have enabled electronic commerce on the Internet. Spectacular growth in commercial use of this new information infrastructure is expected in 1995 and beyond because of the relative ease with which the Mosaic/WWW combination allows for highly accessible information servers to be established on the global network. Because of the decentralized nature of the Internet, it will be difficult to gauge the total business income generated by the introduction of Mosaic; however, there is already enough commercial activity to believe that there will be significant payback on the critical federal HPCCI investment in the NSF supercomputer centers that led to this unexpected software development. Even more important is the paradigm shift in the use of the Internet that Mosaic and the WWW have generated.

BOX 2.2 Solutions to Complex Problems

Parallel computing has enabled creative solutions to a number of industrial and scientific problems. The following examples are but a few of many possible illustrations of such successes.

- **Defense.** Simulation of the interaction between the electromagnetic spectrum and various aircraft design features has enhanced the performance of stealth aircraft. Parallel computing enabled rapid calculations for many different wave lengths and aircraft parts.
- **Petroleum.** Oil exploration and production have been made more productive by three-dimensional analysis and visualization of huge amounts of seismic data. Parallel computing enabled the move from two- to three-dimensional processes.
- **Finance.** Forecasting and simulation of various trading strategies for mortgage-backed security instruments has created a new market and contributed to a reduction in rate spreads. Parallel computing enables simultaneous calculations for numerous instruments within the very short time frames of a fast-moving market. Growth of the finance market arena will provide market pull that should help lower the costs of high-performance processing systems for all types of users.

- The HPCCI has provided large numbers of academic scientists with peer-reviewed access to their choice of high-performance computing architecture to enable their computational science projects. The NSF supercomputing centers, whose core budgets are entirely within the HPCCI, provided access to 23 high-performance computing machines in FY 1995 to 7,500 academic researchers in over 200 universities. The capacity of the centers' supercomputers was 75 times as great as in FY 1986, their first full year of operation, and users represented a broad range of scientific and engineering disciplines (Appendix E lists representative projects).
- The Internet, the flagship of the NREN component of the HPCCI, has become the basis for computer-mediated communication and collaboration among researchers and others. The geographical dispersion of Grand Challenge team members has resulted in pioneering use of electronic collaboration methods, beginning with conventional electronic mail and expanding to multimedia electronic mail, audio and video conferencing, and shared tools for accessing and using remotely stored data and for controlling remote instruments. Development of these methods and tools has been fostered and funded by the HPCCI, demonstrating the potential for electronic collaborators and other approaches to using information technology to support distributed work (CSTB, 1993, 1994e).
- The Internet plus a collection of advances and applications in data storage, analysis, retrieval, and representation—some involving high-performance technology—has catalyzed exploration of digital library concepts. Early Internet-based collaborations among computer scientists, information scientists, cognitive and social scientists, and domain-specific groups provided the basis for a multidisciplinary research effort in digital libraries under the IIITA program that was launched in mid-1993 (NSF, 1993).
- The gigabit network testbeds, an element of the NREN component of the HPCCI, pioneered in advancing the frontiers of communication bandwidth essential to achieving enhancements envisioned for the nation's information infrastructure. In the process, they helped to bridge the gap in perspective and emphasis between the computing and communications research communities.

Transfer of Expertise and Technology

In addition to enabling explicit collaborations, the HPCCI has indirectly affected industry and other sectors outside of academia by stimulating the spread of human experts and thus the transfer of technology, building on a tradition of interaction typical of the computing field. As Kenneth Flamm observed, "People are clearly the medium in which computer technology has been stored and transmitted. The history of the computer industry is the history of trained individuals acquiring knowledge—formal and informal—[and] then applying that knowledge, often in a new organizational or institutional setting" (Flamm, 1988, p. 3).³⁰ The importance of educated and trained talent is reflected in the HPCCI's Basic Research and Human Resources component, which produces the intellectual and human capital, the most general benefit and therefore perhaps the least easy to identify.

Impact on Mission Agencies

In addition to its broad national accomplishments, the HPCCI's contributions to the federal mission agencies, the initial customers for high-performance computing and communications technology, must also be considered. Based on its discussions with agency officials and its own insights into the fit between these technologies and agency activities, the committee believes that the existing and potential contribution of high-performance computing and communications technology to federal mission agencies does justify the investment. The policy decision to eliminate nuclear weapons testing has greatly increased the need for high-performance computer simulations at the Departments of Energy and Defense, for example, and the need to control costs for defense materiel makes simulation in the manufacture of defense-related products an attractive prospect.³¹ Thus much of the HPCCI effort at the Advanced Research Projects Agency (ARPA) relates to design and simulation. Although defense-specific applications are sometimes unique, applications-related investment can have impacts beyond meeting agency needs. As in the case of the gigabit testbeds, for example, such work can provide proofs of concept that encourage private investment by lowering risks.

Five Gigabit Testbed Projects: Collaboration and Impact

The gigabit testbeds provide a case study of how to achieve progress through cross-sectoral, developer-user collaboration to advance high-performance computing and communications technologies. Since 1989, the testbeds have provided the means to test applications and thus extend the state of the art in gigabit networking to link high-performance computers to each other and to applications sites. The five original gigabit testbed projects, funded by NSF and ARPA and administered by the Corporation for National Research Initiatives, were started as a 5-year program.³² They received considerable support from the telecommunications industry, mainly in the form of donated transmission facilities. All three major long-distance carriers and all of the major local-exchange carriers participated. In addition, at least three similar independent testbeds were created through public-private partnerships in the United States, and imitators sprang up in Europe. The testbeds were intended to address two key questions: (1) How does one build a gigabit network?, and (2) What is the utility of such a network?

All five became operational by 1993, showing that gigabit networks could be built. They were largely but not completely successful in illuminating how best to build a gigabit network.³³ The difficulties in many cases were not with the networks but rather the computers connected by the networks, which could not handle the very high bandwidths. Although no research on computer

architecture was included in the testbed projects, the projects demonstrated the need for better computer systems, both hardware and software, to achieve better communications—they demonstrated the intimate linkage between computing and communications systems.

As to the utility of the gigabit testbeds, opinions in the community differ sharply. Demonstrations of the potential utility of gigabit networks in some Grand Challenge applications were achieved, including global climate modeling and radiation treatment planning. What is debated was whether the very high bandwidths actually added much value to the applications. No large-scale use of the testbeds for applications research was really possible because of the rather experimental nature of the networking and limited reach of the networks. Also, this work demonstrated that there were essentially no computers that could take advantage of the high-bandwidth gigabit lines because their internal buses were too slow. Although the gigabit testbeds emphasized speed, discussions within the research community, industry, and the HPCCI agencies have suggested that future testbeds should address architecture, scale, and heterogeneity of systems as well as communications speed.

EVOLUTION OF HPCCI GOALS AND OBJECTIVES

Since early 1994, the policy context for the HPCCI has shifted at least twice, and the change in the Congress heralded by the fall 1994 elections suggests the potential for further change.

Improving the Information Infrastructure

The first shift in policy affecting the HPCCI reflected growing interest in the information infrastructure and thus in the universal, integrated, and automated application of computing and communications technologies to meet a wide variety of user needs. The increasing linkage between the HPCCI and information infrastructure can be seen in the "Blue Books," the principal public documentation of the purpose, scope, participation, budget, achievements, and prospects of the HPCCI.³⁴ Box A.2 in Appendix A outlines the evolution of HPCCI goals as articulated in the Blue Books. Box A.3 indicates the broadening of focus from science to other kinds of applications and drivers of high-performance computing and communications technologies. As discussed above, PL 102-194 marked the first explicit congressional step toward greater attention to information infrastructure; the Blue Books track concurrent thinking of HPCCI agency officials.

The President's FY 1995 budget request bundled the HPCCI together with other items, notably nonresearch programs intended to broaden access to communications and information infrastructure, into the National Information Infrastructure (NII) initiative.³⁵ The initiative built on the level of technology generally available in the early 1990s, proofs of concept provided by the NREN program, and industry trends, including growing use of computer-based networking within and between a variety of organizations and the rise of a variety of network-based businesses. Once the policy focus—in the government, the press, and most of the agencies—centered on information infrastructure, high-performance computing seemed to be greatly downplayed. In many 1993 and 1994 government documents and administration speeches, the first "C" of HPCCI effectively disappeared, notwithstanding the fact that achieving many of the goals for improving the information infrastructure would depend on rapid continuation of progress in high-performance computing. The NII, previously absent, was featured in subtitles of the 1994 and 1995 Blue Books even though there is no formal, specific NII research program.

The formulation of the NII initiative raised questions about the nature and extent of political support for the original HPCCI objectives, and it may have led to expectations that were not embodied in the HPCCI as originally proposed. It perhaps inadvertently underscored the

misperception that because the HPCCI emphasizes the high end of the technology spectrum, it is less relevant or useful than the NII initiative. Cast in more populist terms, the NII initiative included a variety of efforts to explore broadening access to increasingly sophisticated computing and communications services and attention to associated practical concerns. These perceptions—and misperceptions—threaten to slow the momentum of the HPCCI at just the time when its potential to support improvement of the information infrastructure is most needed.

The missing link appears to be the failure of some HPCCI critics to appreciate the dynamism of computing and communications technology: almost by definition, relatively few really need and/or can afford leading-edge computing and communications. But as demonstrated in Chapter 1, the rapid pace of technical development quickly brings these technologies into the mainstream, and they become accessible to a broad populace. Attention to performance is justified by the expectation for rapid transitions from leading-edge technologies to cost-effective, ubiquitous technologies—as well as the kinds of applications expected to grow. For example, multimedia communications will require high-bandwidth, low-delay delivery based on high peak network capacity and on protocol support for negotiating and enforcing service guarantees. The Internet and efforts associated with the development of digital libraries already illustrate the importance of high-performance computing and communications to a broad set of information infrastructure capabilities.

Greater attention to information infrastructure does not imply that performance should be abandoned. But rather than drive toward a narrow goal, such as a teraflop machine or gigabit network, *per se*, the goal should be systems that scale over several orders of magnitude. This goal should include not only processing rates and communication rates, but also storage capacity, data transfer rates, and retrieval times, as well as the problems inherent in serving millions of users.

One can view information technology as a tent: the height of the center pole corresponds to speed and the breadth of the base corresponds to scale (Figure 2.1). Both speed and scale are important research issues. The HPCCI's focus has been mainly, though not exclusively, on speed. We can move toward an enhanced national information infrastructure by adding more cloth to the tent so as to further emphasize scale without deemphasizing speed, or by shifting the focus somewhat from height to breadth, from the research issues of speed to those of scale. Both changes are appropriate; both dimensions are important for the tent to work. Additional opportunities and needs are also suggested by the tent metaphor, recognizing that there is more to advancing information technology and the information infrastructure than speed and scale. Other important goals include:

- Reliability (Will the tent stay up?);
- Software productivity (How long to move the tent to a new site?);
- Malleability (Can the tent's shape be changed?);
- Human-computer interface (Can people use the tent?); and
- Intelligent search and retrieval (Can people find what they want in the tent?).

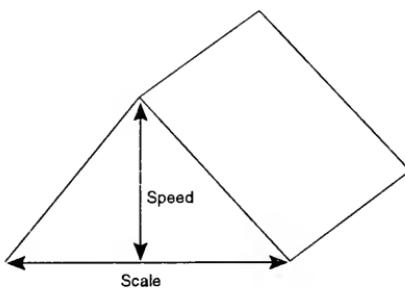


FIGURE 2.1 Scale and speed—important dimensions of the information technology “tent.”

Advancing the information infrastructure presents many practical and some urgent needs, but it has been and can continue to be a driver for the long-term research challenges addressed by the HPCCI. High-performance computing and communications will help provide the technologies

needed to provide flexible, high-rate, affordable data communications. In Office of Management and Budget guidance for developing FY 1996 R&D budgets, the HPCCI is acknowledged as "helping develop the technological foundation upon which the NII will be built," as a prelude to the articulation of several priorities under the broad goal of "harnessing information technology," one of six broad goals.³⁶ See Box 2.3 for an illustrative discussion of how telemedicine needs, for example, can help to drive high-performance computing and communications technology development and deployment, and how the HPCCI can foster paradigm shifts in application domains.

Evolving Research Directions and Relevance for the Information Infrastructure

The public debate over information infrastructure is at heart a debate over how to make computing and communications systems easier to use, more effective, and more productive. The challenge for research policy is to translate usability needs into research topics and programs. The HPCCI itself was built on the recognition that the fundamental challenge to greater acceptance and use of high-performance technologies is to make them more usable. Since the 1970s it has been recognized that more usable parallel processing machines imply the development of algorithms, programming support software, and native parallel applications, but the problem persists despite considerable progress. (See Appendix A.) For information infrastructure in the fullest sense—reaching to ordinary citizens—these efforts must be extended to address intuitive models of use and supporting user interface technologies to enable a class of information appliances that will become a part of everyday life. The acceptance and popularity of Mosaic demonstrate the importance of user models, human factors, and other areas where research is critically needed.

More generally, intelligent information retrieval systems, systems for understanding speech and pictures, and systems for enabling intelligent dialogues between people and computer systems are capabilities that will build on HPCCI research and enhance the usefulness and level of use of information infrastructure. In addition, research and development of core software technologies are needed to achieve progress in security, privacy, network measurement and management, transaction processing, application integration, and other capabilities that may be less directly visible to individuals but that make computing and communications facilities more usable. For example, HPCCI and other computing and communications research can enhance capabilities for distributed, remote measurements of quantities that relate to the environment, traffic flows and management, or health conditions. Yet other research should build on the movement to digital transmission of more and more information. As this list of possibilities suggests, information infrastructure is bigger than an initiative, although one or more initiatives, including the HPCCI, can help to organize and accelerate progress in developing and using it.

Complicating decision making regarding information infrastructure research is the recognition that an advanced information infrastructure is not something that will spring full-grown from any one development. Rather, it is something that will grow from new capabilities in many different sectors of the economy and society. Having to provide for migration, evolution, integration, and interoperability compounds the technical challenges.

**BOX 2.3 Telemedicine: An Example of
HPCCI-Enabled Tele-expertise**

The provision of expert professional services, such as medicine, law, and education, is a current consumer of HPCC technologies as well as a driver of future developments. Generically, this provision of services is often referred to as tele-expertise and can be thought of as the live, interactive provision of services and education between individuals who are geographically separated but electronically connected. Tele-expertise holds the promise of reducing costs and lessening geographical disparity in the availability of services. In particular, telemedicine will be an important part of the National Challenge in health care as evidenced by funding from the National Institutes of Health and other organizations for several projects, including a 3-year contract to use advanced network technologies to deliver health services in West Virginia.

Functionally, telemedicine supplies an audio, visual, and data link designed to maximize understanding between provider and patient. In telemedicine, visual contact and scrutiny are particularly important to accurate communication: studies have suggested that body language and facial expression can convey up to 80 percent of meaning. Clinically, although touch is currently denied, video zoom capability often augments visual examinations beyond what is the norm in face-to-face services. In addition, various endoscopes and physiometers may be utilized across a network to further enhance a health care worker's observations.

Limited telemedicine field trials began in 1958 and expanded with federally funded research demonstrations between 1971 and 1980. Considerable research was done on reliability, diagnostic replicability, user satisfaction, and multiple-specialty services. Currently, a few projects address tele-expertise more broadly by combining telemedicine and distance learning, and trials are being conducted in Montana and Texas that encourage the integrated use of remote services in medicine, industry, law, and education, the "MILES" concept. More specifically, telemedicine has made some advances in the years since the early trials:

- Elaboration and extension of transmission media from early microwave and satellite channels to 384-Kbps service and direct fiber-optic links;
- Reduction of costs due to digital signal compression and decreased long-distance rates in constant dollars; and
- Expansion of the number of pieces of medical equipment that may be connected to the remote terminal, chiefly a variety of endoscopes and physiometers.

Nonetheless, because of health care cost issues and large disparities in the medical services available in different geographical areas, telemedicine has great potential impact as a National Challenge application for HPCCI technologies. Telemedicine urgently needs several HPCCI-related technologies that can be deployed rapidly and inexpensively and that scale well. Among others, these include:

- *Rapid, high-capacity, multipoint switching.* The telephone became increasingly useful as improved switching and networks enabled rapid expansion across the nation. So it is with interactive video—improved switching and networks will activate the distance-spanning benefits of the interactive video market.
- *Translators to interconnect divergent computing and communications technologies.* New technologies are being developed and deployed so rapidly and in so many different places on the globe that it may be more feasible to develop facile, high-performance translators than to struggle for standards.
- *Compact video storage and good retrieval techniques.* Transparent technologies must be developed to enable a physician to efficiently store and easily retrieve salient clinical moments without distracting from the clinical challenge at hand.

SOURCE: Committee on Information and Communication (1994), p. 28.

Although the U.S. telecommunications industry is a world leader in developing and deploying networks on a large scale, the concepts inherent in an advanced national information infrastructure go beyond connecting a large number of relatively homogeneous end systems supporting a relatively small number of applications, such as today's telephones or televisions. Learning how to build large-scale systems, like learning how to build high-performance ones, requires research; it is not simply a matter of deploying lots of boxes and wires. Envisioned for an advanced national information infrastructure is the interconnection of a much larger number and variety of networks than are interconnected today, with more nodes and switches per network and new mixes of wireline and wireless transmission. The end systems of such networks will run a much wider set of applications and call for a broader set of software-based support capabilities often referred to as "middleware." There will be great complexity, increasing the emphasis on scale and architecture and on areas such as accommodating heterogeneity of systems, decentralization of control and management, routing, security, information navigation and filtering, and so on, all of which will depend on software.

The evolutionary nature of information infrastructure also underscores the importance of engaging industry in the planning, support, and conduct of research. Advisory committees and collaborative projects are but two examples of how this engagement can be achieved. See Appendix B, Box B.1, for a discussion of the development of asynchronous transfer mode as an illustration of fruitful industry-university-government interaction.

There have been many government, academic, and industry efforts, some still under way, to identify and clarify research issues associated with improving information infrastructure. The recent CSTB report *Realizing the Information Future* (1994d) provides a unifying conceptual framework from which it derives strategic emphases for research; a multiparty effort generated several lists of research ideas (Vernon et al., 1994); a more focused workshop generated ideas for funding under the NSF networking and communications research program (NSF, 1994); and ARPA's NETS program and other programs have continued to develop and enrich a technology base for bitways and mid-level services to support defense-relevant applications.³⁷

Common to these various efforts is the need for research to enhance such critical information infrastructure middleware capabilities as security and reliability; the basic research underlying many of these concepts had been done by high-performance computing and communications researchers funded mainly by ARPA. In addition, it is important to advance true communications research, including such fundamental areas as transmission, switching, coding and channel access protocols realized in electronic, optical, and wireless technologies, as well as basic computer networking research in such areas as internetworking protocols, transport protocols, flow and congestion control, and so on. These complement and enable efforts relating to distributed computing, which tends to be concerned with the upper or applications level of a total system. See Box 2.4 in the section "Coordination Versus Management" below and Appendix B for an examination of HPCCI communications research efforts and Appendix C for the larger budget allocation picture. Now is the time to explore a wide variety of technical problems, enlisting as many approaches and perspectives as possible.

Overall Computing and Communications R&D Planning

The second major influence on the policy context for HPCCI is a broad rethinking of computing and communications R&D, building on the reorganization of the federal coordinating structures for R&D and factoring in a broad range of technology and policy objectives. The broadest coordination of computing and communications research and development activities across federal agencies is the responsibility of the Committee on Information and Communications (CIC) under the National Science and Technology Council. The CIC was formed in 1994 and is led by

the director of defense research and engineering, the associate director for technology of OSTP, and the assistant director of Computer and Information Science and Engineering, NSF. In late-1994, the CIC launched a strategic planning activity to provide input into the FY 1996 budget-setting process, expected to conclude in early 1995, and inform efforts for the next 5 years. Indications from briefings based on preliminary versions of that strategic plan show a broader and richer set of concerns than previously evident.

Strategic focus areas identified in preliminary materials include global-scale information infrastructure technologies, scalable systems technologies, high-confidence systems, virtual environments, user-centered interfaces and tools, and human resources and education. The HPCCI relates at least somewhat to all of these topics, and the planning process is examining where other, mission-related agency activities can build on HPCCI as well as other activities. Key research activities are classified as components, communications, computing systems, software toolkits, intelligent systems, information management, and applications.³⁸ Software and usability are cross-cutting themes.

Toward a Better Balance

There is a natural evolution of the HPCCI, many aspects of which are associated with improvement of the information infrastructure. The newest component of the HPCCI, the Information Infrastructure Technology and Applications (IITA) program, is one of the most visible signs of this evolution, but also important are the trends within the programs at both ARPA and NSF, which show increasing emphasis on software solutions and tools. ARPA, for example, is devoting attention to software and tools to support design and simulation for development of defense systems; its emphases on security and scalable systems both involve substantial effort relating to software.³⁹ This evolution should continue and indeed accelerate.

Practical experience with the HPCCI and the volatile policy context both suggest that the ideal research agenda for high-performance computing and communications should be driven by strategic priorities, but focused more broadly than on just those priorities. A stable yet flexible approach would combine substantial focus on goals of current national importance, including directly targeted research, with a flexible program that sustains a healthy base of computing, computation, and communications science and technology. The comprehensiveness of the emerging CIC strategic plan appears to provide a broader platform than previously available for supporting the nation's public computing and communications R&D, including that relating to high-performance technology. Also, the commendable inclusion of a technology "maturity model" in CIC's preliminary strategic planning material illustrates recognition of the technology "trickle-down" phenomenon.

MOVING FORWARD—BASIC ISSUES

Balance of Private and Public Investment

The possibility of reduction or even premature termination of the HPCCI, suggested by congressional requests for inquiries by the General Accounting Office (GAO) and the Congressional Budget Office (CBO) and for this committee's report, is troubling. (See Appendix A for a brief discussion of issues raised by GAO and CBO.)⁴⁰ Some HPCCI critics expect industry to pick up the task. They seem to assume possible a larger program of basic research from industry than is

reasonable based on history, the growth in competition, which reduces the profit margins needed to sustain R&D, and the economics of innovation generally.

Leading-edge high-performance computing and communications technology is aimed at the most demanding customers, a niche or subset on the order of 10 percent of the larger computing and communications market. Truly high-end systems tend to be nonstandard and to require considerable customer support, for example, which limits their market potential. It may be more appropriate, therefore, to assume that truly high-end systems are aimed at particular classes of problem for which the systems and associated software have particular value, rather than to assume that these systems will become universal. For example, better weather prediction would save an enormous amount of money and should be carried out on high-performance computers even if millions of people do not have them. The lower end of the market will grow as parallel processing vendors reposition their products, addressing broader industrial and commercial needs for information storage and analysis, multimedia servers, data mining, and intelligent transactions systems.⁴¹

Observers within the computing and communications research communities, including members of this committee, are concerned about the impact of computer and communications industrial restructuring. Changes in the organization of these industries, plus the inherent difficulties incumbent companies face in using research results, prevent companies from undertaking the kind of large-scale, long-range research needed to tackle the challenges inherent in advancing the HPCCI objectives or the broader objectives associated with information infrastructure. This concern is almost impossible to substantiate, because it is inherently intuitive, albeit shaped by expert judgment and the experience of committee members working in or with a variety of computing and communications companies, and because the results of current trends will not be evident for several years.⁴²

Coordination Versus Management

The HPCCI became an integrated, multiagency, cross-cutting initiative because agency and congressional officials recognized that there would be economies of scale and scope from connecting complementary efforts across research disciplines and funding agencies.⁴³ By cooperating, agency officials have successfully leveraged the dollars available in the initiative budget, facilitating collaborative efforts with industry. The NREN infrastructure investments, including the NSFNET backbone and gigabit testbeds, provide examples. Network connections, research tools, and delivery of educational products appear to motivate the broadest interagency activity within the HPCCI context, helping to extend collaborations beyond the conduct of research per se and into a wider circle of agencies.

Through its accomplishments and *esprit de corps*, the HPCCI has become a model for multiagency collaboration.⁴⁴ Each agency retains responsibility for its own part of the program, focusing its participation to meet agency needs and resources. The voluntary compliance of HPCCI agencies with the spirit of PL 102-194 reflects the special cooperation that has characterized the HPCCI. These conditions have enabled the initiative to grow and adapt relatively quickly to changing national needs, technology prospects, and the fit between the two. Perhaps because they see themselves as principal architects of the program, officials from the four initial HPCCI agencies (Department of Defense (DOD), Department of Energy (DOE), National Aeronautics and Space Administration (NASA), and National Science Foundation), in particular, have carried high levels of enthusiasm, dedication, inventiveness, and energy into undertaking the HPCCI. These intangible qualities are widely recognized within the computing and communications research community.

The level of interagency coordination observed today took time to grow. As one might expect when organizations with different missions, budgets, and cultures are faced with a joint task, the HPCCI agencies have disagreed on issues of emphasis and approach over the years. For

example, DOE and NSF have had different views on evolving the NREN program with respect to scope and speed. What is important for the future of the HPCCI, however, is not that differences have arisen but rather that legitimate differences owing to varying missions have been respected, and cooperation and coordination have improved over time. For example, NSF and ARPA—which respectively emphasize small science and larger projects—have worked well in the management of their joint network and computing research activities, as described in Box 2.4.

**BOX 2.4 Coordination in Practice:
The Case of Communications R&D**

The HPCCI currently includes a relatively modest but vigorous communications research program. Three large programs account for \$77 million of the HPCCI communications research budget. Research is concentrated mainly in four areas (see Appendix B for more details and context):

1. Optical networks (the longest-term research),
2. Gigabit networking (medium-term research),
3. Multimedia communications (fairly near term research), and
4. Internetwork scaling (near- and medium-term applied research).

The ARPA networking program, at \$43.1 million, is the largest communications research program activity. The milestones include:

- Demonstration of diverse Internet capabilities such as cable and wireless bitways,
- Demonstration of rate-adaptive quality of service negotiation in asymmetric networks,
- Demonstration of bandwidth and service reservation guarantees for networks in support of real-time services,
- Demonstration of secure routing systems, and
- Interconnection of gigabit testbeds.

The ARPA Global Grid program, at \$23 million, intends to accomplish (in 1995):

- Demonstration of multi-wavelength reconfigurable optical network architecture, and
- Demonstration of integrated DOD and commercial networks in support of crisis management applications.

NSF's Very High Speed Networks and Optical Systems program, at \$11 million, supports research in a wide variety of high-performance networking technologies, including:

- Gigabit testbed research (switching, protocols, and management);
- Resource discovery;
- Information theory;
- Network security;
- Modulation, detection, and coding for information storage; and
- Optical networking.

ARPA and NSF have coordinated well to avoid duplication of efforts. ARPA funds most of the research on internetworking, and NSF concentrates on the deployment of internetworking infrastructure via its NSFNET activities. NSF and ARPA have jointly funded the gigabit testbed research program, which involves demonstration of cross-country gigabit networking technologies.

With the reinforcement provided by PL 102-194, the set of agencies involved in the HPCCI grew. This broader participation better positions the HPCCI to support development and application of computing and communications technologies essential to improving the nation's information infrastructure. However, as reflected in both executive and congressional efforts to promote such

improvements, the information infrastructure raises issues such as deployment, regulation, and other practical aspects that require engaging a broader and somewhat different set of agencies, such as the Federal Communications Commission, the National Telecommunications and Information Administration, and so on, to address a wider range of issues than those relating to R&D.

The diversity of the HPCCI approach allows many views to compete, first for funding, later in the evolution of thinking among researchers, and finally in the marketplace. It also fosters pursuit of the intellectual questions posed by the HPCCI via a range of complementary modes including classical single principal-investigator (PI) research, multiple-PI experimental research, multiple-PI/multiple-field collaborations, intramural research in institutes and national laboratories, and joint industry-government-academia experiments or proofs of concept.

A variety of mechanisms are used to foster interagency cooperation and coordination:

- Joint funding of projects, from relatively specific or narrow activities to the federal portions of the Internet;
- Consortia, such as the consortium on Advanced Modeling of Regional Air Quality involving six federal agencies plus several state and local governments;
- The MetaCenter concept pursued by NSF supercomputer centers (see Appendix E) and extending to other entities and users via the MetaCenter Regional Affiliates program; and
- Cross-agency reviews of grants and contracts, such as the NSF, ARPA, and NASA digital library initiative, and joint testbeds.

The diversity in approach and tactics makes it less likely that the nation will miss some important approach. It also facilitates participation by a variety of agencies, which tend to have different styles as well as emphases for supporting research or procuring technology, consistent with their different missions, histories, and cultures.

As to diversity of mechanisms, the multiple-PI/multiple field category is epitomized by the Grand Challenge teams, which involve multiple institutions attacking frontier research problems with multiple-year horizons, often drawing on access to the leading-edge machines in the NSF supercomputer centers and benefiting from interactions between computer scientists and computational scientists.⁴⁵ The joint industry-government-academia experiment category is currently epitomized by the gigabit network testbeds. More specifically, NASA's FY 1995 HPCCI effort includes integrated multidisciplinary computational aerospace vehicle design and multidisciplinary modeling and analysis of earth and space science phenomena (Holcomb, 1994).

Coordinating Structure

The coordinating structure of the HPCCI has evolved steadily, largely in response to external pressures for improved visibility of decision making, requirements for accountability for expenditures, and the flow of information into and out of the initiative. Some HPCCI observers have continued to argue for a more uniform approach to related activities with thorough planning, precise milestones, and presumably no wasted effort, in a more centralized program. This is the essence of early criticism lodged by the Computer Systems Policy Project (1991 and 1993).

Drawbacks of Centralization

The central question about coordination is whether the special vitality of HPCCI would survive and whether centralized control would convey sufficient benefits, or merely disrupt current arrangements. A more centralized approach would have several drawbacks that could vitiate the

HPCCI: potential loss of variety in perspectives on applications now arising from agencies with different missions; greater risk of concentrating on the wrong thing; and increased bureaucratic overhead and costs associated with efforts to overlay separate programs. Moreover, because much of the existing effort involves previously existing programs, there is a risk that agencies would not participate in a program that involved a loss of their control to a more centralized authority. This concern, probably paramount to the agencies, arises from the recognition that much of the funding associated with the HPCCI is not new, just classified as relevant to the initiative. A virtue of the current arrangement is that the central coordination is provided by a relatively small entity that lacks the resources for micromanagement. That approach maximizes the benefits provided by a diverse group of agencies.

National Coordination Office

BOX 2.5 National Coordination Office: Staffing and Structure through 1994

Chaired by the director of the National Coordination Office (NCO), the High Performance Computing, Communications, and Information Technology (HPCCIT) Subcommittee and its Executive Committee coordinate planning, budgeting, implementation, and program review for the overall initiative. The HPCCIT Subcommittee has also been the major vehicle for communication with other federal agencies, the U.S. Congress, and numerous representatives from the private sector.

The director of the NCO reports to the director of the Office of Science and Technology Policy (OSTP). The director of OSTP has specified that overall budget oversight for the HPCCI be provided by the National Science and Technology Council through the Committee on Information and Communication (CIC). Actual appropriations for the initiative are made in the regular appropriation bills for each participating agency. The HPCCIT coordinates program and budget planning for the initiative and reports to the CIC.

Under the umbrella of the HPCCIT, several working groups have been formed to help guide and coordinate activities within the five components of the initiative. For example, the Information Infrastructure Technology and Applications Task Group, established in 1993, has encouraged and coordinated participating agencies' plans for research and development aimed at providing needed technologies and capabilities for an enhanced nationwide information infrastructure and the National Challenge applications. The Science and Engineering Computing Working Group coordinates activities and software development relating to Grand Challenge applications.

The direct operating expenses of the NCO are jointly borne by the participating agencies in proportion to their HPCCI budgets, and further support is provided by the interagency detailing of staff to the NCO for varying periods of time. Currently the NCO has eight permanent staff and two staff "on loan" from the Department of Energy. In addition to providing general administrative functions such as payroll and personnel administration, the National Library of Medicine also contributes specialized assistance such as public information functions, budget preparation, legislative analysis and tracking, graphic arts services, procurement, and computing and communications support.

SOURCES: See Lindberg (1994); NCO (1994); and CIC (1994). Additional information from letter dated August 8, 1994, to Marjory Blumenthal (CSTB) from Donald A.B. Lindberg (NCO/NLM) in response to committee's interim report (CSTB, 1994c).

The HPCCI coordination focus lies in the National Coordination Office for High-Performance Computing and Communications, which was established in September 1992. It operates under the aegis of the Office of Science and Technology Policy and the National Science and Technology Council (NSTC; see Figure 2.2). Box 2.5 provides information on the structure and staffing of the NCO.

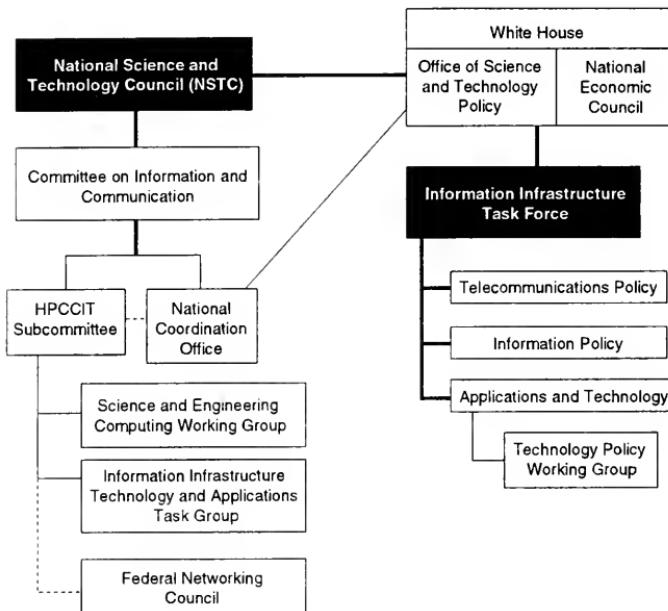


FIGURE 2.2 Organizational context for HPCCI coordination.

The NCO was established to aid interagency cooperation and to serve as liaison for the initiative to the Congress, other levels of government, universities, industry, and the public. It assists the mission agencies in coordinating their separate programs, offering a forum through which the separate agencies can learn of each other's needs, plans, and actions. As part of its coordinating function, the NCO gathers information about the HPCCI activities of different agencies and helps to make this information available to Congress, industry, and the public. Since its formation, the NCO has produced the impressive FY 1994 and FY 1995 Blue Books as visible manifestations of its coordination efforts. The FY 1995 Blue Book is the best documentation available of HPCCI activities and where the money is invested.

Strengthening the NCO. The public debate over the HPCCI attests to the need for improved communication regarding the initiative's purpose and accomplishments. Because lack of external understanding is damaging—not least because it leads to criticisms and investigations that divert energy and resources from pursuing HPCCI goals—the committee believes that the HPCCI

could benefit from a stronger NCO that can do a better job of telling the program's many constituencies about its goals and successes; see the committee's Interim Report (CSTB, 1994c) and Chapter 3.

As authorized in the 1991 High Performance Computing Act, the NCO was to have been assisted by an advisory committee that could provide regular infusions of ideas from industry and academia. To date, the HPCCI has been led mostly by computing visionaries and by people active in science and science applications. That is the right kind of leadership to drive the creation of enabling technology and to create computer architectures that are appropriate for the pursuit of science objectives. But the initiative now also needs the perspective on applications and on making computing and communications technologies more usable that would be provided by an advisory committee of recognized experts with membership balanced between academia and industry, and balanced with respect to application areas and the core technologies underlying the HPCCI.⁴⁶ The growing dependence of more and more people on infrastructure, the rise in potential liabilities of varying kinds, and growth in competitive challenges from abroad increase industry's stakes in the quality of information technology available. Industry input into such issues as standards, security, reliability, and accounting, for example, becomes more important as advancing the information infrastructure and "high-confidence systems" come to drive more of the research agenda.

In lieu of having an advisory committee, the NCO has taken the initiative to convene some industry and other groups to obtain focused input on HPCCI-related issues and directions. In conjunction with its regular meetings with federal HPCCI agency representatives, the NCO has engaged in dialogues with representatives of the computer systems, software, and telecommunications industries; managers of academic computing centers; and others; and it has held a similar discussion with representatives of the mass information storage industry. It has also participated in workshops, conferences, and public meetings sponsored by participating agencies and the subcommittee on High Performance Computing, Communications, and Information Technology (HPCCIT). However, NCO leadership notes that given the restrictions on external interactions imposed by the Federal Advisory Committee Act, the absence of an official advisory committee prevents it from obtaining needed input on an ongoing basis, limiting it instead to one-time interactions and thus foregoing the insights that can arise where parties benefit from repeated interactions.

The committee echoed the NCO's concern by recommending in its interim report that the long-awaited HPCCI advisory committee be established immediately. In view of the delays and difficulties in establishing an HPCCI advisory committee and the apparent tendency of federal science policy leaders to enfold HPCCI in a larger NII initiative, there is some expectation that one advisory committee may be empaneled to provide input into the broader CIC agenda, and to the HPCCI. This National Research Council committee thinks that solution might work, but it urges some action now.

Understanding the Changing Management Context. Actions taken to reinforce the NCO must account for the larger, evolving administrative and management context in which the NCO fits. A key component of that context is the CIC. It receives limited staff support from OSTP and, presumably, from agency-based staff members. Its members include directors of computer- and communications-related research units from across the federal government. Their participation in the CIC provides information exchange and coordination, but the CIC is not an implementation entity. The NCO director participates in both CIC and HPCCIT.

Planning, coordination, and management for the HPCCI have been further confounded by the rise of additional bodies to address technology policy and other policy relating to the NII initiative. The Information Infrastructure Task Force (IITF), formed in 1993, has a Technology Policy Working Group with overlapping representation with the HPCCIT. Its focus is supposed to

be technology policy, versus the HPCCIT's focus on research and development. The IITF receives input from the NII Advisory Committee, also formed in 1993. And on networking issues there are yet more special coordination and advisory entities, such as the Federal Networking Council (FNC) that has associated with it the FNC Advisory Committee. This proliferation of cross-agency entities itself presents many possibilities for confusion. Moreover, by all accounts—from virtually every HPCCI official the committee has heard from and from several private-sector parties—the processes of communication and decision making have been slowed by a calendar-filling profusion of meetings. This situation raises basic practical questions of what work can get done, when, how, and by whom, when committee meetings appear to be the order of the day.

Budget

A detailed overview of the HPCCI budget is presented in Appendix C. The overall level has been subject to misunderstanding.

According to the Blue Books, the HPCCI budget has grown from \$489.4 million in FY 1992 to the \$1.1 billion requested for FY 1995. When the HPCCI was proposed in the executive budget for FY 1992, the agencies involved identified from their existing FY 1991 activities a base that contributed to the goals of the program. The HPCCI's multiagency budget is more complex than it would be had the program been started "from scratch" within a single agency. Although complexity is inherent in multiagency programs and budgets, it has added to the confusion about spending priorities and accomplishments for the initiative.

The agencies that had activities included in the FY 1992 base were the (Defense) ARPA, DOE, NASA, NSF, National Institute of Standards and Technology, National Oceanic and Atmospheric Administration (NOAA), Environmental Protection Agency, and National Institutes of Health/National Library of Medicine. In each subsequent year, agencies have added to this base in two ways: (1) by identifying additional existing programs that contribute to HPCCI goals and (2) by reprogramming and relabeling agency funds to support relevant aspects of the HPCCI. To this base of "identified" activities, Congress has added some funding each year for new activities or the expansion of existing efforts.

The result is that the \$1.1 billion requested for FY 1995 is composed of three elements: (1) funds for the continuation of agency activities that were in existence when the HPCCI started and were designated in the FY 1992 base budget, (2) funds for existing or redirected programs that have since been designated as being a part of the HPCCI, and (3) additional funds for new activities or expansion of existing efforts. It is difficult to determine exactly how large each element is and to make interagency comparisons, because each agency has used slightly different approaches for identifying existing efforts and somewhat different formats for supporting program and budgetary detail. Also, this situation has allowed some agencies (e.g., NOAA) to be considered participants in the initiative without receiving any new money.⁴⁷

NOTES

1. The substance of these components is outlined in the HPCCI's annual Blue Books; see FCCSET (1991, 1993, and 1994).

2. Note that many of the concerns raised here were expressed or articulated in a CSTB report released at the dawn of the HPCCI, *The National Challenge in Computer Science and Technology* (CSTB, 1988).

3. A teraflop refers to 10^{12} (or 1 trillion) floating point operations per second (FLOPS), 1,000 times the performance of the best machines available when the HPCCI began.

4. See Flamm (1988); this book discusses the major computer development projects of the 1940s, 1950s, and 1960s and their dependence on government stimulus and combined government, university, and industry development of technology.

5. In turn, the use of multiple microprocessors in large-scale parallel machines also exposed problems that would have to be resolved for microprocessors as the dominant computing element.

6. See Zatner (1994), pp. 21-25, for a discussion of events leading to Chapter 11 status for TMC. KSR suffered from an accounting scandal, then had to contend with 12 class-action shareholder lawsuits (Snell, 1994). The impact of the lack of software has also been implicated as an indicator of management ineffectiveness in the fates of TMC and KSR (Lewis, 1994).

7. Indeed, talk of the next hurdle, the petaflop system, has already begun. NSF, NASA, DOE, and DOD hosted a 1994 workshop on enabling technologies for petaflop computing. The report is said to argue that that goal can be met at reasonable cost in 20 years using today's paradigms. See Anties (1994), p. 121.

8. "Recommendation A-2: At the apex of the HPC pyramid is a need for a national capability at the highest level of computing power the industry can support with both efficient software and hardware.

"A reasonable goal for the next 2-3 years would be the design, development, and realization of a national teraflop-class capability, subject to the effective implementation of Recommendation B-1 and the development of effective software and computational tools for such a large machine. Such a capability would provide a significant stimulus to commercial development of a prototype high-end commercial HPC system of the future." (NSF, 1993, p. 11)

9. The fundamental computer unit is the microprocessor, which today has a peak speed of around 300 megaflops. It seems premature to build a 3,000 processor teraflop machine in 1995, but as the microprocessors increase in speed to 1 to 2 gigaflops by the late 1990s, it seems reasonable that 512- to 1024-processor teraflop machines may be built if the economics of users and their applications require it. For example, Kenneth Klewer, director of the Center for Computational Sciences at Oak Ridge National Laboratory, was quoted in December 1994 as saying: "The scale here is clearly a function of time, but we could have nearly a teraflop computer today by coupling the Oak Ridge and Los Alamos computers with the ones from Cornell and Maui" (Rowell, 1994). In November 1994, a new product announcement by Japan's NEC indicated that the maximum configuration, with a total of 512 processors, could be rated at a theoretical peak of 1 teraflop (Parker-Smith, 1994b).

The committee notes that if trends at the NSF supercomputer centers continue, the MetaCenter (which pools some of the centers' resources) could achieve an aggregate teraflop in mid-FY 1998 and each center would reach a peak teraflop machine by the end of FY 1999. By contrast, even the aggregate performance would not reach a teraflop until after the year 2000 if acquisition of higher-performance architectures were to revert to pre-HPCCI levels.

10. The HPCCI also triggered considerable debate about what broad availability means—what capabilities, in what locations, accessible by whom and at what cost—anticipating the more recent debates about how universal service in telecommunications should evolve.

11. The gigabit goal, as defined in the NSF-ARPA-CNRI testbeds, was to achieve an end-to-end speed of at least 1 Gbps between two computers on a network. The telephone achievement was to multiplex about 25,000 64-Kbps voice conversations onto a transmission line operating at 1.7 Gbps (late 1980s technology, now more than doubled). The gigabit testbeds have demonstrated end-to-end speeds between two computers of about 500 Mbps, limited by the internal bus speeds of the computers, not the network.

12. Comments by Sandy MacDonald, NOAA, at "Environmental Simulation: The Teraflop and Information Superhighway Imperative Workshop," August 18-20, 1994, Monterey, Calif. He noted an increase from 3,200 numerical observations per day for Kansas in 1985 to 86,000 daily observations, many from automated instruments.

13. Comments by Steve Hammond, National Center for Atmospheric Research, at "Environmental Simulation: The Teraflop and Information Superhighway Imperative Workshop," August 18-20, 1994, Monterey, Calif. He observed that teraflop computing helped reduce processing times to 90 seconds per modeled month, yielding 1,000 modeled months in 30 hours of processing time.

14. Briefings to the committee. Legislative codification and appropriation for a broader vision for the HPCCI have been attempted but have been unsuccessful, most recently in connection with S4/H1757, in 1994.

15. For example, in addressing networking, PL 102-194 also anticipated many of the practical concerns associated with enhancements and expansion of the nation's information infrastructure, such as user charging and protection of intellectual property rights.

16. Policy documents emerging from the administration in mid-1994 and congressional actions in 1993-1994 emphasize a commitment to linking R&D spending to strategic, national concerns (Panetta, 1994; NSTC, 1994b).

17. Of course, there will also be people using the technology who are not in close contact with developers and vice versa—hence the value of a solid base of funding for both computing and communications research and for the sciences that increasingly depend on computation.

18. See CSTB (1988) for a discussion of national challenges within computing. See also CSTB (1992) for a discussion of methods of combining intrinsic problems with problems inspired by needs in other areas.

19. The lengthy time scales associated with developing complex computer-based systems are outlined in CSTB (1994a).

20. Rigorous cost-accounting and auditing can be elaborate, costly, and inflexible: "As a result, R&D done under federal contract is inherently more expensive and less effective than R&D done by an organization using its own funds" (Cohen and Noll, 1994, p. 74).

21. The National Coordination Office (NCO) has taken a first step in the evaluation area through its development of the HPCC implementation plan. The format of the project and program records in that volume provides a basis for subsequent efforts to assess progress. The NCO management has expressed some interest in tracking progress relative to plan elements.

22. The NSF infrastructure program was motivated by the recognition in the late 1970s of major deficiencies in the academic environment for experimental computer science and engineering. The initial stimulus was the report by Feldman and Sutherland (1979). When that report was written, the discipline of computer science and engineering was perceived to be in crisis: faculty members were underpaid relative to research positions in industry and were leaving universities at an increasing rate; the number of new Ph.D.s fell far short of meeting the national demand; most departments lacked experimental computing research facilities; and there was a significant gap in research capability between the top three or four departments, which had benefited from a decade of ARPA investment, and the rest.

23. The Feldman and Sutherland (1979) report resulted in the establishment in 1981 of the Coordinated Experimental Research (CER) program at NSF. The CER program made awards of approximately \$1 million per year (including an institutional match of typically 25 percent) for durations of 5 years to support significant experimental research efforts by building the necessary infrastructure. There have been very substantial increases in the number of departments producing strong experimental research, the number of departments producing strong students in experimental areas, the number of departments conducting leading-edge research in a significant number of areas, the overall rate of Ph.D. production in the field, and other similar measures.

The success of the CER program was important in shaping several subsequent NSF programs that also contributed to the infrastructure of the field, such as the Engineering Research Centers (ERCs) program. A number of ERCs are in computing-related areas, which in turn influenced the Science and Technology Centers (STCs) program; three STCs are in computing-related areas. The CER program itself ultimately became the Research Infrastructure program and was complemented by an Educational Infrastructure program. A number of other agencies instituted research and/or educational infrastructure programs.

24. The NSFNET backbone has involved NSF spending authorized on the order of \$30 million but complemented by in-kind and other investments by IBM and MCI through Advanced Networks and Services, which has deployed and operated NSFNET under a cooperative agreement with NSF. The Internet overall has been growing through proliferation of a variety of commercial Internet access providers. See CSTB (1994d).

25. "Although the infrastructure, including networking, software applications and tools, visualization capabilities, etc., is still not strong enough, raw computing power is becoming comparable, and in some cases greater than what is available at NSF Centers in the U.S. This increase in resources comes at a time when the Japanese government is also increasing its emphasis on basic research for its own needs and to insure that Japan is viewed as a [sic] equitable contributor to the global science community. Readers might want to reflect on the impact the NSF centers have had on U.S. science output and the potential for this to occur in Japan." (Kahaner, 1994b)

26. For example, as a result of detailed interactions between a high-performance computing and communications vendor and a staff member of an NSF supercomputer center, a Grand Challenge computer code uncovered previously undiscovered hardware bugs in newly released microprocessors installed in a scalable supercomputer at the center. This led to the vendor using a version of the Grand Challenge code inside the company as a standard test to uncover both hardware and compiler bugs.

27. Jeremiah P. Ostriker, Princeton University Observatory, personal communication, December 23, 1994.

28. Jeremiah P. Ostriker, Princeton University Observatory, personal communication, December 23, 1994.

29. Examples include the Silicon Graphics Everest/Challenge systems (some 3500 Challenge, Power Challenge, Onyx, and Power Onyx systems were sold in the 15 months following their September 1993 introduction) and the IBM SP2 and Power Parallel systems; see Parker-Smith, 1994a. See also Appendix A.

30. Even the rise and fall of individual ventures shows this generally positive pattern: TMC was launched in part by the expertise of Danny Hillis, previously at MIT, and his associates. With TMC's contraction in 1994, Hillis' team of over 20 engineers from TMC's Future Systems Group went to Sun Microsystems, where they are working on a scalable massively parallel processing system, while other TMC talent continued with a TMC parallel software descendant (Riggs, 1994).

31. Briefings to committee by Victor Reis (Department of Defense) and Howard Frank (Advanced Research Projects Agency).

32. No follow-on program to the gigabit testbed projects has yet been announced. In July 1994, an NSF and ARPA workshop proposed a research agenda for gigabit networking and called for an experimental gigabit research network facility. NSF and ARPA are extending the existing program by a few months, into early 1995.

33. Note that there is continuing popular confusion over the term "gigabit networks" and the fact that the speed most often quoted for them is 640 megabits per second. Each gigabit connection consists of two one-way circuits, each operating at 640 Mbps. Thus the overall speed of the two-way connection is 1.28 gigabits per second when properly compared to the quoted two-way capacity of application networks. Also, the 640-Mbps circuits in at least one case (Aurora) were derived by splitting 2.4-Gbps trunk circuits.

34. Each year beginning in 1991 the director of the Office of Science and Technology Policy submits a report on the HPCCI to accompany the president's budget. The FY 1992, FY 1993, and FY 1994 books were produced by the now-defunct FCCSET; the FY 1995 report was produced by the NCO (acting for the CIC). The report describes prior accomplishments and the future funding and activities for the coming fiscal year. These reports have collectively become known as "Blue Books" after the color of their cover.

35. The NIH initiative was framed in 1993 and included in the FY 1995 budget request.

36. Other goals, such as "a healthy, educated citizenry," also include applications of computing and communications among their priorities.

37. CSTB (1994b); Vemon et al. (1994); and NSF (1994). The ARPA NETS program is covered by the Blue Books for FY 1994 and FY 1995.

38. Briefing to committee by Edward Lazowska, based on a Computing Research Association briefing by John Toole, and augmented by briefings by Anita K. Jones and Howard Frank, December 20, 1994.

39. Briefing to committee by Howard Frank, December 20, 1994.

40. GAO (1993); CBO (1993).

41. Committee briefings by Forest Baskett, Silicon Graphics Inc., April 13, 1991; Justin Ratner, Intel Supercomputer Systems Division, June 27, 1994; Steve Nelson, Cray Research, Inc., June 28, 1994; and Steven Wallach, Convex Computer Corporation, June 28, 1994. Also, see Lewis (1994) re "gigaflops on a budget." See also Furth (1994) for a description of how Encore, Hewlett-Packard, IBM, Pyramid, Tandem, Stratus, and AT&T have changed their focus to transaction processing and fault-tolerant computing.

42. For example, in FY 1994, the NSF centers had an income derived by recovering cycle costs from noncomputer industrial partners of around \$1 million to \$2 million. In comparison to their NSF Cooperative Agreement level of \$16 million per year, this has a small impact. Indeed, the situation is even worse, since a typical NSF supercomputer center receives only half its annual budget from the NSF Cooperation Agreement, the other half coming from state and university matching funds, other grants, and equipment donations by computer vendors.

The center experience also shows that over the last few years, industry spending to attain center know-how—training, software development, information infrastructure application development, virtual reality and visualization projects, and so on—and to use centers as vehicles for collaborative research has increased and exceeds spending on computer processing cycles at some centers. Because the centers have an existing staff for these projects, the industrial income generally covers only the marginal cost of providing that service and therefore does not increase net "new dollars."

43. Other cross-cutting initiatives contemporaneous with HPCCI include advanced manufacturing technology; global change; advanced materials and processing; biotechnology; and science, mathematics, engineering, and technology education (FCCSET, 1993).

44. The HPCCI is understood by a variety of federal officials to have been a model for the "virtual agency" concept advanced through the National Performance Review efforts to improve the organization and effectiveness of the federal government (Gore, 1993).

45. There are 16 grants, 7 awarded in FY 1992 and 9 in FY 1993. Their source of funds can be broken into three parts (NSF/CISE, NSF/non-CISE, and ARPA). The FY 1994 and FY 1995 numbers are shown below. As the chart indicates, CISE's percentage is less than one-third of the funding. This shows great leverage, even greater than that of the centers, roughly one-half of whose budget comes CISE.

	NSF/CISE	NSF/non-CISE	ARPA	TOTAL
FY 1994 \$M	2.77	5.00	1.91	9.68
FY 1994 %	29	52	20	100
FY 1995 \$M	2.79	4.93	1.44	9.16
FY 1995 %	30	54	16	100

46. For example, NASA feels pressure from the HPCCI objectives to orient its program in certain directions, but is encouraged by the aeronautics industry to orient its activities in other directions. NASA is caught in the middle. Aeronautics industry representation at the HPCCI leadership level could help guide the HPCCI in directions that better support the goal of enhancing U.S. industrial competitiveness.

47. Letter dated August 25, 1994 to Marjory Blumenthal from Jerry D. Mahlman (NOAA) in response to committee's interim report (CSTB, 1994c).

Recommendations

The committee believes that strong public support for a broadly based research program in information technology is vital to maintaining U.S. leadership in information technology. Facilitated access for both academic and industrial users to advanced computing and communications technologies has produced further benefits both in scientific progress and in U.S. industrial competitiveness. The committee's recommendations for the High Performance Computing and Communications Initiative (HPCCI) are based on this view of the importance of information technology to the country, as well as on the track record of success for the government's investment in information technology research. The committee's 13 recommendations address five different areas:

- General research program;
- High-performance computing;
- Networking and information infrastructure, including work focusing on the National Challenges;
- The supercomputer centers and the Grand Challenge projects; and
- Coordination and program management.

Within each area the recommendations are presented in priority order.

GENERAL RECOMMENDATIONS

As discussed in Chapter 1, government investment has played a major role in maintaining U.S. leadership in information technology and in helping to advance the technology, providing benefits to virtually every citizen. The return on federal investment has been substantial.

Recommendation 1. Continue to support research in information technology. Ensure that the major funding agencies, especially the National Science Foundation and the Advanced Research Projects Agency, have strong programs for computing research that are independent of any special initiatives. Past investment has yielded significant returns, as demonstrated in Chapter 1. Continued government investment in computing research, *at least as high as the current dollar level*, is critical to continuing the innovation essential to maintaining U.S. leadership.

Today the HPCCI supports nearly all of this research, an arrangement that is both misleading and dangerous: misleading because much important computing research addresses areas other than high performance (even though it may legitimately fit under the new Information Infrastructure Technology and Applications (IITA) component of the HPCCI), and dangerous

because reduced funding for the HPCCI could cripple all of computing research. The "war on cancer" did not support all of biomedical research, and neither should the HPCCI or any future initiative on the nation's information infrastructure subsume all of computing research.

Recommendation 2. Continue the HPCCI, maintaining today's increased emphasis on the research challenges posed by the nation's evolving information infrastructure. In addition to the work on infrastructure carried out in the new IITA program, continuing progress is needed in areas addressed by the HPCCI's other four components (High-Performance Computing Systems, National Research and Education Network (NREN), Advanced Software Technology and Algorithms, and Basic Research and Human Resources).

The NSFNET and the gigabit testbeds have demonstrated the ability to build larger-scale, higher-performance networks, but ongoing research in several areas is still needed before a ubiquitous high-performance information infrastructure can be developed and deployed nationwide. The committee supports the HPCCI's increasing focus on information infrastructure, emphasizing that successful evolution of the nation's communications capability rests on continued investment in basic hardware, networking, and software technologies research. To further this evolution, which is consistent with administration efforts, including the addition of the IITA program, plus General Accounting Office (GAO, 1994) and other recommendations, the committee has identified in Recommendations 3 through 10 program areas that should receive (a) increased emphasis, (b) stay at present levels, and (c) have reduced federal support.

RECOMMENDATIONS ON HIGH-PERFORMANCE COMPUTING

Recommendation 3. Continue funding a strong experimental research program in software and algorithms for parallel machines. It is widely recognized that software for parallel computers lags behind hardware development. Progress in software and algorithms for parallel computers will determine how quickly and how easily we can use them.¹ A shift in emphasis toward increased funding for software and algorithm activities under high-performance computing has already begun. This shift properly reflects the urgency of investing more in software.

The committee recommends the following approach to continue progress in research areas critical to developing and building needed software and algorithms:

- Continue research on compilers, programming languages, and tools aimed at making it easier to use parallel computing machines. Critical needs include improved portability across machines, improved ability to run programs on machines of different sizes, and better understanding of how best to use different multiprocessor memory organizations.
- Continue to develop experimental operating systems for parallel computers. More operating system experience will help us learn how to improve parallel hardware. Focus on the underlying research challenges posed by parallel machines rather than developing commercial operating systems technology.
- Continue research on database and information systems for parallel machines. Such applications have increased in importance and represent a promising area for using parallel computing.
- Continue research in the use of parallel computing for graphics and visualization. Graphics applications are valuable both because they demand much from their software and hardware and because they stimulate effective use of high-performance computing by

offering computational scientists and other end-users the ability to analyze complex data and problems.

- Fund sufficient hardware purchases to ensure needed access for computer scientists and end users trying to evaluate the effectiveness of new architectures and software technologies. Dedicated access to expensive machines is often required for operating systems development or for controlled measurement of software performance, and sometimes dedicated access is needed to full-scale machines, which are then most economically housed in a centralized, national center. The importance of the local availability of mid-range machines for researchers in software for parallel computers was noted in the Branscomb report (NSF, 1993).
- Provide resources to help complete the development and distribution of compilers, programming tools, and related infrastructure broadly usable by the software research community. Such infrastructure—which may be developed by individual research groups or by centers (such as the NSF science and technology centers)—has been crucial to rapid progress. For example, tools for the design of very large scale integrated (VLSI) circuits allowed many researchers to undertake VLSI designs. The committee notes that funding agencies should avoid turning related infrastructure development efforts into product development efforts.
- Seek improved integration of parallel computing hardware and software with communication networks, both in software and hardware research.
- Emphasize design and analysis of new algorithms for parallel computing, as well as implementation and evaluation of these algorithms on real parallel machines. Opportunities for development of new parallel algorithms exist in both scientific and information infrastructure-related applications. The theoretical performance and scaling efficiency of new algorithms need to be demonstrated by actual implementation and evaluation on parallel machines, first by computer scientists and then embedded in real end-user applications.
- Ensure that effective new algorithms for parallel computing are made widely available to end-user communities to assist in building applications.

Recommendation 3.1. Avoid funding the transfer (“porting”) of existing commercial applications to new parallel computing machines unless there is a specific research need.

Several existing applications enjoy widespread commercial use on large uniprocessor and vector machines; examples include third-party codes in chemistry, biomolecules, engineering fluid dynamics, deformable structures, and database access. It has been proposed by some that the HPCCI should fund transferring, or “porting,” such applications to new types of parallel computers as a way to enhance the attractiveness of new parallel machines. The committee finds inappropriate the use of federal HPCCI funding for such porting of applications for several reasons. First, the algorithms used in these applications were designed for sequential or vector computing, and thus little new knowledge will be gained from merely porting existing applications to a parallel machine without redesigning the algorithms. Second, the open market will fund such transfers if a sufficient user base exists. Third, choosing whose application to transfer and to which machines will involve the HPCCI in picking winners from among many commercial vendors.

Although it recognizes that a federal agency might decide that one of its missions would best be served by porting an existing application to a parallel computer, the committee recommends that funding of such ports be justified on the basis of the agency mission and not as HPCCI

research. The committee believes that it is legitimate for groups of agencies to work together to develop community codes for common applications needed by their several missions. Likewise, carrying out an HPCCI research program may require that applications be available on a particular parallel machine, in which case the transfer could be justified by the importance of the research it enabled. Finally, the committee also sees a legitimate reason to port existing applications for the purpose of evaluating machines within a research laboratory or university.

Recommendation 4. Stop direct HPCCI funding for development of commercial hardware by computer vendors and for "industrial stimulus" purchases of hardware. Maintain HPCCI support for precompetitive research in computer architecture; this work should be done in universities or in university-industry collaborations and should be driven by the needs of system and application software. The development and placement of parallel hardware to date was necessary to establish parallel computing as a viable alternative to sequential and vector computing. (The establishment of this paradigm is discussed throughout Chapter 2 and in Appendices A and E.) Industry is now willing and able to improve on the base of ideas established by the HPCCI, at least for mainstream parallel machines (special government requirements are discussed below). Government development funds should no longer be spent in industry either to further refine parallel machines or to purchase machines as a stimulus for vendors.

The committee notes that use of HPCCI funds for these purposes has already decreased significantly, a trend that the committee supports.² Federal funding of hardware developments within companies should continue to decline, unless some special agency need demands the development of nonstandard hardware (e.g., a high-performance system for use on a ship or in an airplane, such as Intel Corporation's parallel Paragon computer made more rugged for military use, or for a highly specialized application). In such cases, agency mission funds, and not HPCCI funds, should be used.

Important precompetitive hardware research problems merit continued federal funding because the development of parallel computing architecture and gigabit networks will not be the final chapter in the continuing development of ever more powerful systems. The committee recommends that ongoing research efforts in hardware and architecture be based in academic and research institutions, possibly in collaboration with industry. Potential problems can be minimized if the research institution serves as the project lead, and if the research challenges rather than commercial development are the focus (Cohen and Noll, 1994). Not only do academic institutions have more freedom to think about longer-term issues, but they also stimulate technology transfer through publication and placement of graduates. The national experience supports a basic tenet of Vannevar Bush: publicly funded research carried out in universities produces excellence, diversity, fresh ideas, trained people, and technology transfer (OSRD and Bush, 1945). Commercial organizations, on the other hand, have powerful incentives to avoid distributing new ideas widely and may even impede the introduction of new technology when it competes with existing products.

To narrow the gap between parallel computing hardware capabilities and the software needed to use them, research on architecture should be driven by software and applications needs. Thus, further integration of application and system software needs into architecture research should be encouraged in any funding of architecture research.

Recommendation 5. Treat development of a teraflop computer as a research direction rather than a destination. The committee believes that federal investment in developing or purchasing machines to demonstrate raw scalability for its own sake is inappropriate, except as a focus for precompetitive, academic research. Instead, the focus should be on matching agencies' mission requirements to the emerging sustainable scalable architectures. Such architectures will very likely reach 1-teraflop capability before the end of this decade using 1,000 or so high-performance commercial microprocessors.

The goals of scalability over many sizes of machine and of demonstrating teraflop performance have been useful in pointing toward the use of mass-produced devices in large

collections to solve complex computing tasks, but implementation of a machine of any specific size can be premature. Moreover, seeking a common design over a large size range is wasteful because the expensive communication paths required to harness large numbers of inexpensive processors together are inappropriate when scaled down to smaller machines with only a few processors. The pursuit of wide scalability may have deferred early consideration of shared-memory parallel computers, the type that today appears promising. In fact, the focus on teraflop capability detracts from other important aspects of high performance, such as memory and input/output systems, which are critical components of any high-performance system.

Advances in parallel architectures together with progress in the underlying integrated circuit technology will continue to provide improvements in performance/cost ratios that will naturally bring computing power to the teraflop level. Most industry analysts see the potential for single microprocessors with 1- to 2-gigaflop peak performance by the end of the decade. Combining 512 to 1024 such future microprocessors in a scalable system would create a teraflop capacity at roughly the price of today's supercomputers, with capabilities of tens of gigaflops possible. Supporting research into the key technologies needed to achieve and use scalable computing, combined with patience to see how the relative economics of computing power and communications interact, seems to this committee to be the most efficient approach to increasing performance.

The committee thus emphasizes that the HPCCI should treat the goal of teraflop performance as a milestone to be reached naturally by computer vendors in due course, not on a forced time scale. The HPCCI should continue to fund research on technologies that will contribute to reaching the goal. At some point in the near future a teraflop parallel machine will be built when some agencies' mission requirements correspond to a sufficiently economical commercial offering. Continued progress will naturally lead to machines much larger than a teraflop.

RECOMMENDATIONS ON NETWORKING AND INFORMATION INFRASTRUCTURE

The committee believes that the successes of the HPCCI in establishing scalable compute servers, investigating high-performance networks, and forming interdisciplinary teams of computer and application scientists are setting the stage for important new research to support enhancement of the nation's information infrastructure. An increased emphasis on the research needed to achieve such an infrastructure is desirable (CSTB, 1994d).

In fact, this shift has already begun; spending on networking and IITA activities accounted for nearly 50 percent of the HPCCI budget requested in FY 1995:³ \$177 million for the NREN and \$282 million for IITA. This is a significant increase over the \$114 million that was spent for NREN in FY 1993, the year prior to the addition of IITA. The committee believes that such a shift is appropriate.

Recommendation 6. Increase the HPCCI focus on communications and networking research, especially on the challenges inherent in scale and physical distribution. Advancing the nation's information infrastructure will put great demands on digital communications technology for providing broad access to services. Ensuring broad access poses a host of technical and economic questions for which existing solutions are inadequate. The committee recommends increased support for learning how to attach millions of users to a digital communications structure that provides a wide array of services and greater integration of services, and how to accommodate the demands that these users will generate using the novel applications enabled by such an information infrastructure.

Recommendation 7. Develop a research program to address the research challenges underlying our ability to build very large, reliable, high-performance, distributed information systems based on the existing HPCCI foundation. An improved infrastructure will need to offer

capability to all facets of our economy on a scale not yet imagined, and no one yet anticipates all of the ways that users will use such an information infrastructure.

Improvements to the nation's information infrastructure and activities related to it have generated a level of public interest matched by only a few technology-based objectives. The committee is concerned that unrealistic expectations for availability and for the quality and range of services could encourage a short-term, product-oriented focus in funding research activities⁴ that would not be in our country's best interest. Care should be taken to apprise policymakers and the public of the long time needed for development and wide-scale deployment of the services expected to be available through the information infrastructure.

The committee strongly recommends that the HPCCI remain focused on the basic research issues arising from desired improvements to the information infrastructure, evolving from its early emphasis on parallel, high-performance computing, high-performance networking, and scientific applications to the broader issues of connection, scale, distributed systems, and applications. The addition of the IIITA area to the HPCCI was a key step in accelerating a shift in focus of the research community to the challenges of improving the nation's information infrastructure. The committee has identified three key areas where new emphasis is critical to supporting the research needs associated with the information infrastructure:

1. *Scalability.* While the HPCCI has emphasized large computing systems on the order of thousands of interacting computers, an enhanced, nationwide information infrastructure will require scaling to millions of users. In addition, the HPCCI has emphasized bringing the highest performance to bear on individual scientific applications, whereas improving the information infrastructure for the nation requires using the highest performance to meet the practical needs of millions of simultaneous users.

2. *Physical distribution and the problems it raises.* A better information infrastructure will emphasize geographical distribution with its limitations on bandwidth, increase in latency of communication, and additional challenges in secure and reliable communication. These challenges have been much less severe in localized high-performance parallel systems. Research on both distributed and parallel systems technology will be important in supporting this aspect of a national-scale information infrastructure.

3. *Innovative applications.* A shift should occur from a focus on specific Grand Challenge problems in science to well-formulated National Challenges that affect a wider segment of society. The committee sees an important role for development and demonstration of easily appreciated applications that will drive the technology of the information infrastructure.

Improving scalability and physical distribution requires investment in both:

- *Hardware and architecture*, including systems that efficiently handle a rich mix of text, images, and audio and video data; systems that provide cost-effective, high-bandwidth, end-to-end communications; and systems that provide information access to large numbers of users; and
- *Software*, including basic networking software for encryption, routing, flow control, and so on; tools for providing and building such capabilities as scheduling, bandwidth optimization, video handling, and service adaptation; and many others. This is the so-called "middleware."

The committee believes that building on the current HPCCI model of a coordinated program, avoiding central control, seems even more crucial for the IITA portion of the research program, because the challenges posed by an information infrastructure are inherently less well defined than those addressed by other components of the initiative. The committee is encouraged by the early development of cooperative research programs in IITA areas, such as the digital libraries program, which includes NSF, ARPA, and the National Aeronautics and Space Administration (NASA), and by recent attempts to identify topics for research, such as discussions among several hundred researchers and others at a workshop in early 1994 (Vernon et al., 1994).

Recommendation 8. Ensure that research programs focusing on the National Challenges contribute to the development of information infrastructure technologies as well as to the development of new applications and paradigms. The National Challenges incorporate socially significant problems of national importance that can also drive the development of information infrastructure. Hardware and software researchers should play a major role in these projects to facilitate progress and to improve the communication with researchers developing basic technologies for the information infrastructure. Awards to address the National Challenges should reflect the importance of the area as well as the research team's strength in both the applications and the underlying technologies. The dual emphasis recommended by the committee contrasts with the narrower focus on scientific results that has driven many of the Grand Challenge projects.

Because the National Challenges as currently defined are too broad and vague to offer specific targets for large-scale research, the notion of establishing testbeds for a complete national challenge is premature. Instead, research funding agencies should regard the National Challenges as general areas from which to select specific projects for limited-scale testbeds or focused software research projects. Particular areas in which a focused research target can be identified (e.g., the ARPA-NSF-NASA digital library testbeds) may be appropriate for slightly higher funding, but the committee believes that very large scale applications development is premature. At this early stage, letting "a thousand flowers bloom" will provide a better pay-back than funding a few large or full-scale deployments. (Box A.3 and related text in Appendix A give more information on the National Challenges.)

RECOMMENDATIONS ON THE SUPERCOMPUTER CENTERS AND GRAND CHALLENGE PROGRAM

The four NSF supercomputer centers are the largest single element of the FY 1995 HPCCI implementation plan in dollars (\$76 million, or 6.6 percent of the requested FY 1995 HPCCI budget) and the largest infrastructure project in the initiative. The centers—which give users access to a broad array of powerful tools ranging from highly innovative to mature—are a significant national resource for gaining knowledge, experience, and capability. Thanks to their leadership, high-performance computing and communications are now widely accepted as an important tool in academia, industry, and commerce.

The centers have played a major role in establishing parallel computing as a full partner with the prior paradigms of scalar and vector computing. They have contributed by providing access to state-of-the-art computing facilities to a broad range of users. As new large-scale architectures appear, the centers stimulate their early use by providing access to these architectures and by educating and training users. (Appendix E details the accomplishments of the NSF centers and of their national user base.)

The committee recognizes that advanced computation is an important tool for scientists and engineers and that support for adequate computer access must be a part of the NSF research program in all disciplines. The committee also sees value in providing large-scale, centralized computing, storage, and visualization resources that can provide unique capabilities. How such

resources should be funded and what the long-term role of the centers should be with respect to both new and maturing computing architectures are critical questions that NSF should reexamine in detail, perhaps via the newly announced Ad Hoc Task Force on the Future of the NSF Supercomputer Centers Program. For example, much of the general access to computing resources at the centers is provided on maturing architectures. Neither the quality of the science performed by the users of such technology nor the appropriateness of NSF funding for such general access is in question. The committee did not consider the appropriate overall funding level for the centers. Nonetheless, the committee does question the exclusive use by the NSF of HPCCI-specific funds for support of general computing access, which in itself does not simultaneously help drive the development of high-performance computing and communications technology.

In this regard, NSF follows a different approach to funding its supercomputer centers than do NASA and the Department of Energy (DOE), where HPCCI funds are used only to support the exploration and use of new computing architectures, while non-HPCCI funds are used to support general access. The committee believes that DOE's and NASA's approach to funding general access should be followed across the agencies. Also, as the committee points out in Recommendation 12, including all of the NSF supercomputer centers' funding under HPCCI could cause major disruption to the centers' national mission should the HPCCI be altered significantly.

Recommendation 9. *The mission of the National Science Foundation supercomputer centers remains important, but the NSF should continue to evaluate new directions, alternative funding mechanisms, new administrative structures, and the overall program level of the centers. NSF could continue funding of the centers at the current level or alter that level, but it should continue using HPCCI funds to support applications that contribute to the evolution of the underlying computing and communications technologies, while support for general access by application scientists to maturing architectures should come increasingly from non-HPCCI funds.*

Examination of the supercomputer centers program should include identification of:

- Emerging new roles for the centers in supporting changing national needs; and
- Future funding mechanisms, including charging mechanisms and funding coupled to disciplinary directorates.

In addition to enabling high-performance scientific computing, several of the NSF centers have developed new software technologies that have significantly affected other parts of the HPCCI. The most obvious of these is the recently developed Mosaic World Wide Web browser. The committee recommends that NSF continue to take a broad view of the centers' mission of providing access to HPCCI resources, including, but clearly not limited to, participating in research needed for improved software for parallel machines and for enhancement of the nation's information infrastructure. The centers, and the researchers who use their facilities, should compete for research funds by the normal means established by the funding agencies.

Recommendation 10. *The Grand Challenge program is an innovative approach to creating interdisciplinary and multi-institutional scientific research teams; however, continued use of HPCCI funds is appropriate only when the research contributes significantly to the development of new high-performance computing and communications hardware or software. Grand Challenge projects funded under the HPCCI should be evaluated on the basis of their contributions both to high-performance computing and communications and to the application area.* The Grand Challenge problems are sufficiently large and complex and the research teams addressing them are capable enough to exercise the parallel computing technology thoroughly and to test its capability. These efforts have been supported under the HPCCI as a valid way to involve real users in parallel computing, but as parallel computing becomes an established tool, the need to use the HPCCI to stimulate the user community will decrease. Furthermore, the use of

high-performance computing will become more pervasive, making it unreasonable to include all such programs with the HPCCI.

The committee recommends completion of the initial Grand Challenges as planned over the next few years. Currently, although the scientific disciplines are providing major funding for Grand Challenge projects (e.g., more than 50 percent of requested FY 1995 funds for NSF Grand Challenges come from the scientific and engineering research directorates), virtually all of the Grand Challenge funding is labeled HPCCI. The committee urges that any follow-on funding of Grand Challenges include a significant and growing fraction of non-HPCCI scientific disciplinary funds. This will limit the selection to tasks whose scientific interest justifies their cost, in competition with other research in their respective disciplines.

The committee sees an ongoing value from the strong interaction between challenging applications and new architectures and software systems and from cooperation between computer and computational scientists—a number of the Grand Challenge teams have demonstrated that collaboration can lead to advances in both computing and the particular scientific discipline involved. Partial funding of applications research that contributes to the development of new hardware and software systems is a legitimate use of HPCCI funds. Such activities must be evaluated on the basis of their contributions both to high-performance computing and communications technologies and to the application area.

RECOMMENDATIONS ON COORDINATION AND PROGRAM MANAGEMENT IN THE HPCCI

Recommendation 11. Strengthen the HPCCI National Coordination Office (NCO) while retaining the cooperative structure of the HPCCI and increasing the opportunity for external input. As the committee pointed out in its interim report (CSTB, 1994c, p. 9), the dimensions of the need for clear communication about the HPCCI have recently become apparent: congressional oversight activities and other indicators suggest that the HPCCI is of concern to a growing constituency and that often a variety of audiences need detailed explanations of it. Such an effort will add substantially to the work of the NCO, which has been headed by a half-time, permanent-position director who holds a concurrent, half-time appointment as director of the National Library of Medicine (NLM).⁵ The other NCO staff positions are a mix of permanent positions, contract positions, and temporary positions filled by individuals on loan from other federal agencies for limited periods of time, often no more than 1 year.⁶

Although the NCO reports to the Office of Science and Technology Policy (OSTP) on programmatic matters, administrative functions such as office space, salaries, and benefits have been handled through the National Institutes of Health. The temporary nature of some of the NCO positions jeopardizes continuity and cumulative insight. Further, limited staff resources raise questions about the NCO's capacity to meet the challenge of the growing volume, complexity, and urgency of the outreach efforts for the initiative (CSTB, 1994c, p. 9).

The NCO serves an important coordination and communication function both among agencies of the government and between the agencies, Congress, industry, and the public. It is to the credit of the NCO staff that the HPCCI has been an effective model of interagency collaboration. In recommending a strengthening of the NCO, the committee strongly endorses the current NCO's role of supporting the mission agencies rather than directing them. The committee believes that it is vital that direction of the HPCCI remain in the agencies.

By avoiding actual direction the NCO leaves mission judgments in the hands of responsible agency officials who are accountable for the allocation of their resources. By avoiding the appearance of direction the NCO encourages an appropriate diversity of research projects as each agency capitalizes on its best ideas. Mission agencies cooperate effectively with each other and

with the NCO precisely because it does not threaten their autonomy. This cooperation could easily vanish were the NCO seen as functioning with too heavy a hand. The committee believes that the value of interagency cooperation outweighs the potential benefits that might be gained through more centralized management of the HPCCI (CSTB, 1994c, p. 8).

The committee strongly recommends retaining the model of a cooperative and coordinated interagency program. Some individuals and organizations have expressed concern about the lack of centralized management of the HPCCI. However, the committee believes that the current cooperative structure is one of the initiative's strengths, providing room for diversity of thought and action. Such diversity is essential to progress, because no one manager or agency has a monopoly on the right ideas for the future of computing and communications. Central management of the HPCCI could focus its activities too narrowly, as well as lead to reduced interest in the program by agencies that found that the centralized agenda no longer matched their interests.⁷

The committee believes that government investment in information technology research has often enjoyed first-class leadership. Program officers with vision have supported innovative ideas, leading to later successes. The committee emphasizes that the best method for making continued research investment is to continue to attract highly competent program officers and to give them the flexibility to develop effective programs. In the past, this approach has yielded solid returns on the research investment. Furthermore, it has encouraged necessary diversity in the research program, thus increasing opportunities for unexpected discoveries and ensuring a broad perspective in addressing problems.

Recommendation 11.1. Immediately appoint the congressionally mandated advisory committee intended to provide broad-based, active input to the HPCCI, or provide an effective alternative. The HPCCI could be improved by input from and review by an advisory committee with balanced representation from industry and academia, including current and potential users and developers of high-performance computing and communications. If appointment of such a committee is not feasible, some alternate mechanism should soon be developed to provide similar input. The committee is aware of the recent increases in the number of advisory committees, as well as the danger of having too many committees. Thus, the committee recommends that the HPCCI advisory committee have a well-defined role focusing primarily on providing external input into the coordination and planning for the HPCCI.

Recommendation 11.2. Appoint an individual to be a full-time coordinator, program spokesperson, and advocate for the HPCCI. Having a part-time NCO director has served well to this point, but the broadening of the HPCCI demands leadership unencumbered by other major responsibilities. A full-time person could either direct the NCO or could report to the director and would work to strengthen the ties between the HPCCI, industry, the Congress, and the public. The committee uses the word "coordinator" rather than "manager" to emphasize the need for coordination and communication that avoid usurping the authority of the mission agencies. The NCO should remain within the OSTP structure.

Recommendation 12. Place projects in the HPCCI only if they match well to its objectives. A number of preexisting agency programs have entered the HPCCI. Such administrative changes make the HPCCI budget appear to grow faster than the real growth of investment in high-performance computing and communications. Some of these programs exactly match the goals of the HPCCI and are properly included. Others are only marginally relevant and might better be placed elsewhere in agency budgets. The committee sees the possibility of a long-term danger to important programs, such as basic research in computing within NSF and ARPA, should the HPCCI end.

Recommendation 12.1. Federal research funding agencies should promptly document the extent to which HPCCI funding is supporting important long-term research areas whose future funding should be independent of the future of the HPCCI. The committee found that many research areas predating the HPCCI and related only partly to its goals are now under the

HPCCI umbrella. Although encouraging important research areas to include and even focus on HPCCI-related components, the process of coding all funding in a research area as high-performance computing and communications can be dangerous. In many cases, areas were recoded as high-performance computing and communications without receiving an increase in funds. The danger in this process is that changes in the direction or level of funding for the HPCCI could lead to unintentional changes in the funding levels of important research areas, even if they are largely unrelated to the HPCCI and even if they have received none of the HPCCI incremental funding!

This problem is particularly acute at NSF, where nearly all of the funding in the Computer and Information Science and Engineering directorate is coded as HPCCI funding. Given that NSF is not a mission agency and is investigator-driven, this approach seems shortsighted. NSF would have to retain funding for computer science research even in the absence of the HPCCI. Ongoing funding of important research areas in computer science will be critical to the nation's future, independent of the future of the HPCCI.

Recommendation 13. Base mission agency computer procurements on mission needs only, and encourage making equipment procurement decisions at the lowest practical management level. To stimulate the use of parallel computing early in the HPCCI's 5-year time frame, it has been appropriate and necessary to place into service a reasonable number of highly parallel machines for serious algorithm and software development. Early development of an adequate base of parallel computers was essential to shifting the attention of industry and research organizations toward parallel computing. Now, however, it is more appropriate to base procurement of computer hardware on mission needs only. One program that claims to have done so already is the Defense High Performance Computing Modernization Program. The committee applauds modernization of the computing facilities available to Department of Defense organizations and the mission-driven nature of the procurement process, which should be established in all agencies.

Each agency has responsibility for its own budget and its own requirements. The committee believes that agencies should base procurements of computing equipment on their needs and budget constraints. Agencies should be free to purchase parallel computers when they suit agency needs. Individual agencies can balance the cost of obtaining applications against the cost of computing equipment so as to best match procurements to their requirements. Parallel computing is now mature enough to be considered a viable alternative to other forms of computing and may deliver suitable computing power at less cost than other architectures.

Although the committee firmly recommends that computer purchases be guided by mission needs, it also sees a role for collaboration between mission agencies and industrial or university parallel computing consortia. Direct agency responsibility for missions, budgets, and equipment purchases can be reconciled with the advantages of group action through participation in appropriate consortia. For example, the NSF centers have proven valuable for offering exploratory experience with high-performance computing, and it is encouraging to see industrial-academic consortia forming to explore parallel computing. The committee encourages mission agencies to participate with the NSF centers and other parallel computing consortia. Such participation offers knowledge at low cost and leads ultimately to more cost-effective procurements.

The committee's recommendation that equipment be selected at the lowest practical management level applies equally to government agencies and to government contractors. The direct manager of a computing facility is charged with making it work and will do that task best if allowed to select equipment that matches the facility's needs. The committee believes that agency-wide procurement of standard brands, while promoting collaboration, can weaken the responsibility of the user organizations. Likewise, it has generally been best for an agency to specify the results it wants and to leave the choice of specific equipment to the contractor.⁸ Delegating equipment selection not only saves top-level agency decision-making resources but also places responsibility for purchase decisions firmly in the hands of the managers who must deliver results.

**COMMENTS RELATING THIS REPORT'S RECOMMENDATIONS FOR
HIGH-PERFORMANCE COMPUTING AND COMMUNICATIONS
RESEARCH TO ADMINISTRATION PRIORITIES**

This report's recommendations broadly address much of the computer science and engineering research being conducted today, as well as the HPCCI specifically. In a May 1994 memorandum from the director of the Office of Management and Budget to all agency heads, the administration outlined its priorities for U.S. research and development in general. Box 3.1 briefly compares applicable parts of that memorandum to the positions taken and actions recommended in this report of the Committee to Study High Performance Computing and Communications.

**BOX 3.1 Comparison of Administration Priorities for
Harnessing Information Technology to Committee
Recommendations in This Report**

Computing Systems

"The development of scalable systems with the input/output capabilities, mass storage systems, real-time services, and information security features needed to build and fully utilize the National Information Infrastructure (NII) should be emphasized. High performance computing systems capable of 10^{12} operations per second (a teraflop) on technical problems will be achieved by FY 1997. Emphasis should be placed on advances in information storage media for both high and low end applications; systems integration of clustered workstations and large parallel systems; development of advanced tools and processes for the design and prototyping of faster semiconductor devices; and research on nanotechnology, photonics, flat panel displays, and integrated micro-electrical-mechanical devices."

The Committee on High Performance Computing and Communications recommends continuing the focus on high-performance parallel computing, but decreasing the emphasis on achieving a teraflop computing system on a specific time scale. The committee agrees that scalable high-performance computing and communications technology that supports the nation's emerging information infrastructure is important. The committee also recommends increased emphasis on the support of communications within new computer systems.

Networking and Communications

"It will be necessary to support the development of the networking technology required for the deployment of national gigabit speed networks incorporating heterogeneous carriers including satellite and wireless capability. This means serving hundreds of millions of users and demonstrating mobile and wireless capability. It includes the development of interoperability concepts and technologies and the integration of computers, televisions, telephones, wireless telecommunications and satellites."

The committee thoroughly agrees.

continues

BOX 3.1—continued***Software, Algorithms, and Basic Research***

"The United States should conduct basic research to support the computational requirements of new computing paradigms. There is a need for new methods for data authentication and software verification and validation. The development of tools and techniques to enable assembly of systems from inexpensive, versatile, reusable software components is required."

The committee's recommendations emphasize the importance of a complete program of basic research in computer science and underscore the importance of research on software and algorithms to make the best use of new computing paradigms.

Information Infrastructure Services

"Access to and utilization of the NII will require services, tools, and interfaces that facilitate a wide range of applications. These include registries, directories, navigation and resource discovery tools, data interchange formats, and other information services that help users find and query services and components in distributed repositories. There will have to be new types of human-computer access and the development of improved collaborative software, groupware, and authoring tools for multimedia will be important. Equally important are the development of privacy and security technologies and integrated information and monitoring systems."

The committee enthusiastically agrees.

Human-Computer Interaction

"New products and applications are enabled by those hardware and software technologies that will allow every American to use easily the NII. Development and use of the following should be advanced: virtual reality; simulation; flat-panel displays; video and high definition systems; three-dimensional sound, speech interfaces, and vision."

The committee did not study human-computer interaction and specific related areas for research, although these are certainly a key part of a broad research program on information infrastructure.

Computing and Communications Applications

"The FY 1996 R&D budget should advance applications of high-performance computing and the NII. The Federal High-performance Computing and Communications Program is helping to develop the technologies and techniques needed to solve critical research problems that require more advanced computers, storage devices, algorithms, and software tools. Additional effort is needed to accelerate the transfer of these technologies from the laboratory to the marketplace."

The committee recommends that applications be funded increasingly from outside the HPCCI, especially as the technology underlying the application becomes mature. HPCCI funds should focus research support on applications that affect the base computing and communication technologies, as well as solve new applications problems. The committee recommends that such a policy be followed in funding work on the Grand Challenges and the emerging National Challenges.

SOURCE: Panetta (1994), pp. 10-11.

NOTES

1. See the subsections "Programming" and "Algorithms" in Appendix A for a discussion of development and achievements to date.
2. Based on briefings to the committee by Anita Jones (Department of Defense), Duane Adams, Howard Frank, and John Toole (Advanced Research Projects Agency); and Victor Reis (Department of Energy).
3. See NCO (1994), p. 15. Note that figures represent the President's requested budget authority for FY 1995. Actual appropriated levels were not available at press time. Because the HPCCI is synthesized as a cross-cutting multiagency initiative, there is no direct "HPCCI appropriation."
4. This risk is illustrated in the GAO examples of standards setting and other nonresearch activities under the HPCCI umbrella.
5. A January 6, 1995, press release from the Office of Science and Technology Policy announced the resignation of the NCO's director, Donald A.B. Lindberg. Lindberg, requesting that a successor be named when his current 2-year term ends, recommended that a full-time director be appointed at this point in the evolution of the HPCCI.
6. Letter dated August 8, 1994, to Marjory Blumenthal (CSTB) from Donald A.B. Lindberg (NCO/NLM) in response to the committee's interim report (CSTB, 1994c).
7. The committee shares OSTP Director John Gibbons' concerns about the centralized management advocated by GAO (1994, p. 34).
8. The committee notes the extreme fruitfulness of this model in the hands of the late Sidney Fernbach of DOE's Lawrence Livermore National Laboratory, who for a generation kept his laboratory at the forefront of computing and at the same time helped to stimulate the development of generations of supercomputers.

Bibliography

- Anthes, Gary H. 1994. "Cray Envisions New Frontier," *Computerworld*, July 18, p. 121.
- Anthes, Gary H., and Mitch Betts. 1994. "R&D: Measure of Success," *Computerworld*, November 14, p. 32.
- Browder, Felix E. 1992. "Of Time, Intelligence, and Institutions," *Dædalus* 121(1 Winter):105-110.
- Carlton, Jim. 1994. "Makers of PCs Had Very, Very Happy Holiday," *Wall Street Journal*, December 27, pp. B1 and B4.
- Carnevale, Mary Lu. 1994. "U.S. Awards Grants to Spur Data Network," *Wall Street Journal*, October 13, p. B6.
- Cohen, Linda R., and Roger G. Noll. 1994. "Privatizing Public Research," *Scientific American* 271(3):72-77.
- Commission on Physical Sciences, Mathematics, and Applications (CPSMA); National Research Council. 1994. *Quantitative Assessments of the Physical and Mathematical Sciences: A Summary of Lessons Learned*. National Academy Press, Washington, D.C.
- Committee on Information and Communication (CIC). 1994. *High Performance Computing and Communications: Technology for the National Information Infrastructure, Supplement to the President's Fiscal Year 1995 Budget*. National Science and Technology Council, Washington, D.C.
- Computer and Business Equipment Manufacturers Association (CBEMA). 1994. *Information Technology Industry Data Book 1960-2004*. Computer and Business Equipment Manufacturers Association, Washington, D.C.
- Computer Science and Technology Board (CSTB), National Research Council. 1988. *The National Challenge in Computer Science and Technology*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1992. *Computing the Future: A Broader Agenda for Computer Science and Engineering*. National Academy Press, Washington, D.C.

- Computer Science and Telecommunications Board (CSTB), National Research Council. 1993. *National Collaboratories: Applying Information Technology for Scientific Research*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1994a. *Academic Careers for Experimental Computer Scientists and Engineers*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1994b. *Information Technology in the Service Society: A Twenty-first Century Lever*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1994c. *Interim Report on the Status of the High Performance Computing and Communications Initiative*. Computer Science and Telecommunications Board, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1994d. *Realizing the Information Future: The Internet and Beyond*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1994e. *Research Recommendations to Facilitate Distributed Work*. National Academy Press, Washington, D.C.
- Computer Science and Telecommunications Board (CSTB), National Research Council. 1995. *Information Technology for Manufacturing: A Research Agenda*. National Academy Press, Washington, D.C., in press.
- Computer Select. 1994. *Data Sources Report*. Ziff Communications Company, New York, December.
- Computer Systems Policy Project (CSPP). 1991. *Expanding the Vision of High Performance Computing and Communications: Linking America for the Future*. Computer Systems Policy Project, Washington, D.C., December 3.
- Computer Systems Policy Project (CSPP). 1993. *Perspectives on the National Information Infrastructure: CSPP's Vision and Recommendations for Action*. Computer Systems Policy Project, Washington, D.C.
- Computer Systems Policy Project (CSPP). 1994. *Perspectives on the National Information Infrastructure: Accelerating the Development and Deployment of NII Technologies*. Computer Systems Policy Project, Washington, D.C.
- Congressional Budget Office (CBO). 1993. *Promoting High-Performance Computing and Communications*. U.S. Government Printing Office, Washington, D.C., June.
- Corcoran, Elizabeth. 1994. "The Changing Role of U.S. Corporate Research Labs," *Research-Technology Management* 37(4):14-20.

- Council on Competitiveness. 1991. *Gaining New Ground: Technology Priorities for America's Future*. Council on Competitiveness, Washington, D.C.
- Coy, Peter. 1993. "R&D Scoreboard: In the Labs, the Fight to Spend Less, Get More," *Business Week*, June 28, pp. 102-124.
- Deng, Yuefan, James Glimm, and David H. Sharp. 1992. "Perspectives on Parallel Computing," *Daedalus* 121(1 Winter):31-52.
- Economides, Nicholas, and C. Himmelberg. 1994. "Critical Mass and Network Size," presented at the Telecommunications Policy Research Conference, Solomons Island, Md., October 1-3.
- Faltermayer, Edmund. 1993. "Invest or Die," *Fortune* 127(4):42-52.
- Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Office of Science and Technology Policy. 1991. *Grand Challenges: High Performance Computing and Communications, The FY 1992 U.S. Research and Development Program*. Committee on Physical, Mathematical, and Engineering Sciences, Office of Science and Technology Policy, Washington, D.C., February 5.
- Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Office of Science and Technology Policy. 1992. *Grand Challenges 1993: High Performance Computing and Communications, The FY 1993 U.S. Research and Development Program*. Committee on Physical, Mathematical, and Engineering Sciences, Office of Science and Technology Policy, Washington, D.C.
- Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Office of Science and Technology Policy. 1993. *FCCSET Initiatives in the FY 1994 Budget*. Office of Science and Technology Policy, Washington, D.C., April 8.
- Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), Office of Science and Technology Policy. 1994. *High Performance Computing and Communications: Toward a National Information Infrastructure*. Committee on Physical, Mathematical, and Engineering Sciences, Office of Science and Technology Policy, Washington, D.C.
- Feldman, Jerome A., and William R. Sutherland. 1979. "Rejuvenating Experimental Computer Science," *Communications of the ACM*, September, pp. 497-502.
- Flamm, Kenneth. 1987. *Targeting the Computer: Government Support and International Competition*. Brookings Institution, Washington, D.C.
- Flamm, Kenneth. 1988. *Creating the Computer: Government, Industry, and High Technology*. Brookings Institution, Washington, D.C.
- Forrester Research Inc. 1994. "Home PCs: The Golden Age," *The Forrester Report*, October.
- Furht, Borko. 1994. "Parallel Computing: Glory and Collapse," *Computer* 27(11):74-75.
- Geppert, Linda. 1994. "Industrial R&D: The New Priorities," *IEEE Spectrum* 31(9):30-41.

- Goldfarb, Debra. 1994. *Workstations and High-Performance Systems: Primary Applications for High-Performance Computers*. International Data Corporation, Framingham, Mass.
- Gore, Jr., Albert. 1993. *From Red Tape to Results, Creating a Government That Works Better & Costs Less: Reengineering Through Information Technology, Accompanying Report of the National Performance Review*. U.S. Government Printing Office, Washington, D.C., September.
- Hamilton, Alexander. 1791. "Report [to the Congress] on Manufactures," in M.J. Frisch (ed.), *Selected Writings and Speeches of Alexander Hamilton*, pp. 277-318. American Enterprise Institute for Public Policy Research, Washington, D.C., and London.
- Hillis, W. Daniel. 1992. "What Is Massively Parallel Computing, and Why Is It Important?" *Dædalus* 121(1 Winter):1-15.
- Holcomb, Lee B. 1994. Statement before the House Committee on Science, Space, and Technology: Subcommittee on Science, May 10.
- Information Infrastructure Task Force. 1993. *The National Information Infrastructure: Agenda for Action*. Information Infrastructure Task Force, Washington, D.C., September 15.
- International Business Machines Corporation. 1954. *Preliminary Report: Specifications for the IBM Mathematical FORMula TRANslating System, FORTRAN*. Programming Research Group, Applied Science Division. International Business Machines, New York.
- International Business Machines Corporation. 1956. *The FORTRAN Automatic Coding System for the IBM 704 EDPM (Programmer's Reference Manual)*. International Business Machines, New York.
- Kahaner, David K. 1994a. "Fujitsu Parallel Computing Workshop, 11/94, Kawasaki," November 28, distributed electronically.
- Kahaner, David. 1994b. "Hitachi Parallel Computer Announcement," April 14, distributed electronically.
- Kahaner, David K. 1994c. "Japanese Supercomputer Purchases," September 2, distributed electronically.
- Kaufmann III, William J., and Larry L. Smarr. 1993. *Supercomputing and the Transformation of Science*. Scientific American Library, New York.
- Lewis, Ted. 1994. "Supercomputers Ain't So Super," *Computer* 27(11):5-6.
- Lindberg, Donald A.B. 1994. Statement before the House Committee on Science, Space, and Technology: Subcommittee on Science, May 10.
- Markoff, John. 1995. "Digital Devices Draw Consumers," *New York Times*, January 7, pp. 39 and 50.

National Coordination Office (NCO) for High Performance Computing and Communications. 1993. *Technology Transfer Activities Within the Federal High Performance Computing and Communications Program*. NCO HPCC 93-01. National Coordination Office for High Performance Computing and Communications, Bethesda, Md., April.

National Coordination Office (NCO) for High Performance Computing and Communications, Office of Science and Technology Policy. 1994. *FY 1995 Implementation Plan*. National Coordination Office for High Performance Computing and Communications, Bethesda, Md.

National Science and Technology Council (NSTC). 1994a. "Program Description." Electronics Subcommittee, Civilian Industrial Technology Committee, National Science and Technology Council, Washington, D.C.

National Science and Technology Council (NSTC). 1994b. *Science in the National Interest*. National Science and Technology Council, Washington, D.C., August.

National Science Foundation (NSF). 1993. *Research on Digital Libraries: Announcement*. NSF 93-141. National Science Foundation, Washington, D.C.

National Science Foundation (NSF). 1994. *Research Priorities in Networking and Communications*. Report to the NSF Division of Networking and Communications Research and Infrastructure by Members of the Workshop held May 12-14, 1994, Airlie House, Va.

Nelson, David. 1994. Statement before the House Committee on Science, Space, and Technology: Subcommittee on Science, May 10.

New York Times. 1994. "Experimental Fusion Reactor at Princeton Sets a Record," November 9, p. A22.

NSF Blue Ribbon Panel on High Performance Computing. 1993. *From Desktop to Teraflop: Exploiting the U.S. Lead in High Performance Computing*. National Science Foundation, Arlington, Va., August.

Office of Science and Technology Policy. 1989. *The Federal High Performance Computing Program*. Executive Office of the President, September 8.

Office of Scientific Research and Development (OSRD) and Vannevar Bush. 1945. *Science: The Endless Frontier*. U.S. Government Printing Office, Washington, D.C.

Panetta, Leon E. 1994. "Memorandum for the Heads of Departments and Agencies: FY 1996 Research and Development (R&D) Priorities." Office of Management and Budget, Washington, D.C., May 12.

Parker-Smith, Norris (editor at large). 1994a. "Hi-demand, Lo-price Dilemma for Scalable Vendors," *HPCwire* 3(45):Article 1518 (distributed electronically).

Parker-Smith, Norris (editor at large). 1994b. "NEC Joins Parallel Parade with Wide Range of CMOS Systems," *HPCwire* 3(45, November 11):Article 1513, distributed electronically.

- Patterson, David A., and John L. Hennessy. 1994. *Computer Organization and Design: The Hardware/Software Interface*. Morgan Kaufmann Publishers, San Mateo, Calif.
- Rensberger, Boyce. 1994. "Scientific Ranks Outpace Funds: Imbalance May Put Nation's Technological Primacy at Risk." *Washington Post*, December 25, pp. A1 and A20.
- Riggs, Brian. 1994. "IBM, Sun Vie for Thinking Machines' MPP Technology," *Computer* 27(12):9.
- Roach, Stephen S. 1994. "Economics Inside the U.S. Economy," *U.S. Investment Research*. Morgan Stanley, New York, July 15.
- Rowell, Jan. 1994. "ORNL's Ken Kiewer on How to Put the Sizzle Back in HPCC," *HPCwire* 3(47, December 2):Article 9023, distributed electronically.
- Schatz, Bruce R., and Joseph B. Hardin. 1994. "NCSA Mosaic and the World Wide Web: Global Hypermedia Protocols for the Internet," *Science* 265(12 August):895-901.
- Schwartz, Jacob T. 1992. "America's Economic-Technological Agenda for the 1990s," *Dædalus* 121(1, Winter):139-165.
- Sikorovsky, Elizabeth. 1994. "ARPA Issues \$85M Networking Bonanza," *Federal Computer Week*, April 25, p. 3.
- * Snell, Monica. 1994. "Supercomputing: Same Name, Different Game," *Computer* 27(11):6-7.
- Toole, John C. 1994. Statement before the House Committee on Science, Space, and Technology: Subcommittee on Science, May 10.
- U.S. Department of Commerce (DOC), International Trade Administration. 1994. *U.S. Industrial Outlook, 1994*. U.S. Government Printing Office, Washington, D.C.
- U.S. Department of Defense. 1960. *COBOL: Initial Specifications for a Common Business Oriented Language*. U.S. Government Printing Office, Washington, D.C.
- U.S. Department of Defense. 1988. *Bolstering Defense Industrial Competitiveness: Report to the Secretary of Defense by the Undersecretary of Defense (Acquisition)*. Department of Defense, Washington, D.C.
- U.S. General Accounting Office (GAO). 1993. *High Performance Computing: Advanced Research Projects Agency Should Do More to Foster Program Goals*. U.S. General Accounting Office, Washington, D.C., May.
- U.S. General Accounting Office (GAO). 1994. *High Performance Computing and Communications: New Program Direction Would Benefit from a More Focused Effort*. GAO/AIMD-95-6. U.S. General Accounting Office, Washington, D.C., November.
- Vernon, Mary K., Edward D. Lazowska, and Stewart D. Personick (eds.). 1994. *R&D for the NII: Technical Challenges*. Report of a workshop held February 28 and March 1, 1994, in Gaithersburg, Md. EDUCOM, Washington, D.C.

Vishkin, Uzi (ed.). 1994. *Developing a Computer Science Agenda for High-Performance Computing*. ACM Press, New York.

Zatner, Aaron. 1994. "Sinking Machines," *Boston Globe*, September 6, pp. 21-25.

A**The High Performance Computing
and Communications Initiative: Background****THE TECHNICAL-ECONOMIC IMPERATIVE
FOR PARALLEL COMPUTING****The United States Needs More
Powerful Computers and Communications**

The High Performance Computing and Communications Initiative (HPCCI) addresses demanding applications in many diverse segments of the nation's economy and society. In information technology, government has often had to solve larger problems earlier than other sectors of society. Government and the rest of society, however, have mostly the same applications, and all find their current applications growing in size, complexity, and mission centrality. All sectors are alike in their demands for continual improvement in computer speed, memory size, communications bandwidth, and large-scale switching. As more power becomes increasingly available and economical, new high-value applications become feasible. In recent decades, for example, inexpensive computer power has enabled magnetic resonance imaging, hurricane prediction, and sophisticated materials design. Box A.1 lists additional selected examples of recent and potential applications of high-performance computing and communications technologies. (See also Appendix D for a list of applications and activities associated with the "National Challenges" and Appendix E for an outline of supercomputing applications.)

**BOX A.1 Examples of Important Applications
of High-Performance Computing and Communications Technologies**

- Continuous, on-line processing of millions of financial transactions
- Understanding of human joint mechanics
- Modeling of blood circulation in the human heart
- Prediction and modeling of severe storms
- Oil reservoir modeling
- Design of aerospace vehicles
- Linking of researchers and science classrooms
- Digital libraries
- Improved access to government information

Conventional Supercomputers Face Cost Barriers

For four decades and over six computer generations, there has been a countable demand, much of it arising from defense needs, for a few score to a few hundred supercomputers, machines built to be as fast as the state of the art would allow. These machines have cost from \$5 million to \$25 million each (*in current* dollars). The small market size has always meant that a large part of the per-machine cost has been development cost, tens to hundreds of millions of dollars. Such products are peculiarly susceptible to cost-rise, market-drop spirals.

As supercomputers have become faster, they have become ever more difficult and costly to design, build, and maintain. Conventional supercomputers use exotic electronic components, many of which have few other uses. Because of the limited supercomputer market, these components are manufactured in small quantities at correspondingly high cost. Increasingly, this cost is capital cost for the special manufacturing processes required, and development cost for pushing the state of the component and circuit art.

Moreover, supercomputers' large central memories require high bandwidth and fast circuits. The speed and complexity of the processors and memories demand special wiring. Supercomputers require expensive cooling systems and consume large amounts of electrical power. Thoughtful prediction shows that supercomputers face nonlinear cost increases for designing and developing entirely new circuits, chip processes, capital equipment, specialized software, and the machines themselves.

At the same time, the end of the Cold War has eliminated much of the historical market for speed at any cost. Many observers believe we are at, or within one machine generation of, the end of the specialized-technology supercomputer line.

Small Computers Are Becoming Faster, Cheaper, and More Widely Used

Meanwhile the opposite cost-volume spiral is occurring in microcomputers. Mass-production of integrated circuits yields single-chip microprocessors of surprising power, particularly in comparison to their cost. The economics of the industry mean that it is less expensive to build more transistors than to build faster transistors. The per-transistor price of each of the millions of transistors in mass-produced microprocessor chips is extremely low, even though their switching speeds are now quite respectable in comparison to those of the very fastest transistors, and a single chip will now hold a quite complex computer.

While microprocessors do not have the memory bandwidth of supercomputers, the 300-megaflop performance of single-chip processors such as the MIPS 8000 is about one-third the 1-gigaflop performance of each processor in the Cray C-90, a very fast supercomputer. Microprocessor development projects costing hundreds of millions of dollars now produce computing chips with millions of transistors each, and these chips can be sold for a few hundred dollars apiece.

Moreover, because of their greater numbers, software development for small machines proves much more profitable than for large machines. Thus an enormous body of software is available for microprocessor-based computers, whereas only limited software is available for supercomputers.

Parallel Computers: High Performance for Radically Lower Cost

Mass-production economics for hardware and software argue insistently for assembling many microcomputers with their cheap memories into high-performance computers, as an alternative to developing specialized high-performance technology. The idea dates from the 1960s, but the confluence of technical and economic forces for doing so has become much more powerful now than ever before.

CHALLENGES OF PARALLEL COMPUTING

Organizing a coherent simultaneous attack on a single problem by many minds has been a major management challenge for centuries. Organizing a coherent simultaneous attack on a single problem by a large number of processors is similarly difficult. This is the fundamental challenge of parallel computing. It has several aspects.

Applications

It is not evident that every application can be subdivided for a parallel attack. Many believe there are classes of applications that are inherently sequential and can never be parallelized. For example, certain phases in the compilation of a program are by nature sequential processes.

Many applications are naturally parallel. Whenever one wants to solve a problem for a large number of independent input datasets, for example, these can be parcelled out among processors very simply. Such problems can be termed intrinsically parallel.

Most applications lie somewhere in between. There are parts that are readily parallelized, and there are parts that seem sequential. The challenge is how to accomplish as much parallelization as is inherently possible. A second challenge of great importance is how to do this *automatically* when one starts with a sequential formulation of the problem solution, perhaps an already existing program.

Hardware Design

How best to connect lots of microprocessors together with each other and with shared resources such as memory and input/output has become a subject of considerable technical exploration and debate. Early attempts to realize the potential performance of parallel processing revealed that too rigid a connection between machines stifles their ability to work by forcing them into lock-step. Too loose a connection makes communication between them cumbersome and slow. The section below, "Parallel Architectures," sketches some of the design approaches that have been pursued.

Numerical Algorithms

During the centuries of hand calculations, people worked one step at a time. Ever since computers were introduced, the programs run on them have been mainly sequential, taking one small step at a time and accomplishing their work rapidly because of the prodigious number of steps that they can take in a short time. The current numerical algorithms for attacking problems are

mostly sequential. Even when the mathematics of solution have allowed high degrees of parallel attack, sequential methods have generally been used. In fact, most languages used to express programs, such as FORTRAN, COBOL, and C, enforce sequential organizations on operations that are not inherently sequential.

In the 30 years since parallel computers were conceived, computational scientists have been researching parallel algorithms and rethinking numerical methods for parallel application. This work proceeded slowly, however, because there were few parallel machines from which to benefit if one did come up with a good parallel algorithm, and few on which to develop and test such an algorithm. People didn't work on parallel algorithms because they had no parallel machines to motivate the work; people didn't buy parallel machines because there were few parallel algorithms to make them pay off. The HPCCI and its predecessor initiatives broke this vicious cycle. By funding the development of machines for which little market was developed, and by providing them to computational scientists to use, the HPCCI has vastly multiplied the research efforts on parallel computation algorithms. Nevertheless, 30 years of work on parallel approaches has not yet caught up with four centuries of work on sequential calculation.

Learning New Modes of Thought

Programmers have always been trained to think sequentially. Thinking about numerous steps taken in parallel instead of sequentially may initially seem unnatural. It often requires partitioning a problem in space as well in time. Parallel programming requires new programming languages that can accept suitable statements of the programmer's intent as well as new patterns of thought not yet understood and formalized, much less routinely taught to programmers.

A NEW PARADIGM

By responding to the technological imperative for parallel computing, the HPCCI has in a major way helped add a new paradigm to computing's quiver. Parallel computing is an additional paradigm, not a replacement for sequential and vector computing. Large numbers of researchers have begun to understand the task of harnessing parallel machines and are debating the merits of different parallel architectures. Because the parallel paradigm is new, no one can yet say which particular approaches will prove most successful. It is clear however, that this healthy debate and the workings of the market will identify and develop the best solutions.

Has the parallel computing paradigm really been established as the proper direction for high-performance computing? The Committee to Study High Performance Computing and Communications: Status of a Major Initiative unanimously believes that it has. It is obliged to report, however, that the issue is still being hotly debated in the technical literature. In the November 1994 special issue of the Institute of Electrical and Electronic Engineers *Computer* magazine, Borko Furht asserts in "Parallel Computing: Glory and Collapse" that "the market for massively parallel computers has collapsed, but [academic] researchers are still in love with parallel computing." Furht (1994) argues, "We should stop developing parallel algorithms and languages. We should stop inventing interconnection networks for massively parallel computers. And we should stop teaching courses on advanced parallel programming." An editorial by Lewis (1994) in the same issue similarly discounts highly parallel computing.

Part of the difference of opinion is semantics. Computers have had a few processors working concurrently, at least internal input/output processors, since the late 1950s. Modern vector supercomputers have typically had four or eight processors. The new paradigm concerns *highly* parallel computing, by which some mean hundreds of processors. The committee believes that the

number of processors in "parallel" computers in the field will grow normally from a few processors, to a few dozen processors, and thence to hundreds. For the next several years, many computer systems will use moderate parallelism.

The strongest evidence, and that which convinces the committee that the parallel computing paradigm is a long-term trend and not just a bubble, comes from the surging sales of third-generation parallel processors such as the SGI Challenge, the SGI Onyx, and the IBM SP-2. SGI's director of marketing reports, for example, that SGI has sold more than 3,500 Challenge and Onyx machines since September 1993; IBM's Wladawsky-Berger reports that 230 orders for the SP2 have been booked since it was announced in summer 1994 (Parker-Smith, 1994a). In fairness to Furht and Lewis, these surging sales figures have appeared only in the last few months, whereas journal lead times are long.

COMPUTER ARCHITECTURES

Overview

Sequential, Vector

The simple sequential computer fetches and executes instructions from memory one after the other. Each instruction performs a single operation, such as adding, multiplying, or storing one piece of data. Decisions are made by conditionally changing the sequence of instructions depending on the result of some comparison or other operation. Every computer includes memory to store data and results, an instruction unit that fetches and interprets the instructions, and an arithmetic unit that performs the operations (see Figure A.1).

Vector computers perform the same instruction on each element of a vector or a pair of vectors. A vector is a set of elements of the same type, such as the numbers in a column of a table. So a single "add" operation in a vector computer can cause, for example, one column of 200 numbers to be added, element by element, to another column of 200 numbers. Vector computers can be faster than sequential computers because they do not have to fetch as many instructions to process a given set of data items. Moreover, because the same operation will be done on each element, the flow of vector elements through the arithmetic unit can be pipelined and the operations overlapped on multiple arithmetic units to get higher performance.

Parallel

Parallel computers also have multiple arithmetic units, intended to operate at the same time, or in parallel, rather than in pipelined fashion. Three basic configurations are distinguished according to how many instruction units there are and according to how units communicate with each other. Within each configuration, designs also differ in the patterns, called topologies, for

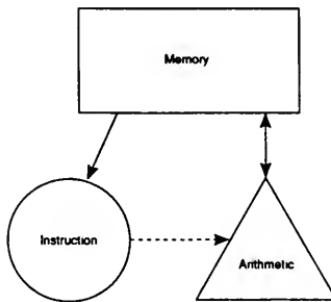


FIGURE A.1 Sequential computer organization.

connecting the units to each other to share computational results. Thus applications programmed for a particular computer are not readily portable, even to other computers with the same basic configuration but different topologies.

Single Instruction Multiple Data. In a single instruction multiple data (SIMD) computer, one instruction unit governs the actions of many arithmetic units, each with its own memory. All the arithmetic units march in lock step, obeying the one instruction unit but fetching different operands from their own memories (Figure A.2). Because of the lock step, if any node has to do extra work because of the particular or exceptional values of its data, all the nodes must wait until uniform operations can proceed. Full utilization of all the processor power depends on highly uniform applications.

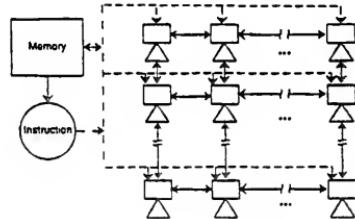


FIGURE A.2 A data-parallel computer organization.

Multiple Instruction Multiple Data

Message Passing. In a multiple instruction multiple data (MIMD) message-passing computer, each arithmetic unit has its own memory and its own instruction unit. So each node of such a machine is a complete sequential computer, and each can operate independently. The multiple nodes are connected by communication channels, which may be ordinary computer networks or which may be faster and more efficient paths if all the nodes are in one cabinet. The several nodes coordinate their work on a common problem by passing messages back and forth to each other (Figure A.3). This message-passing takes time and instructions. Various topologies are used to accelerate message routing, which can get complex and take many cycles.

There are two quite different forms of MIMD computers, distinguished by the network interconnecting the processors. One, commonly called a massively parallel processor (MPP), has a collection of processor nodes co-located inside a common cabinet with very high performance specialized interconnections.

The other, often called a workstation farm, consists of a group of workstations connected by conventional local area or even wide area networks. Such collections have demonstrated considerable success on large computing problems that require only modest internode traffic. Between the two extremes of the MPP and the workstation farm lie a number of parallel architectures now being explored. No one can say how this exploration will come out.

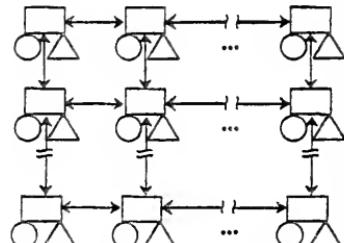


FIGURE A.3 A message-passing parallel computer organization.

Multiple Instruction Multiple Data Shared-Memory. In a multiple instruction multiple data (MIMD) shared-memory computer, the separate nodes share a large common memory. The several nodes coordinate work on a common problem by changing variables in the shared memory, which is a simple and fast operation (Figure A.4).

Each node also has its own memory, generally organized as a *cache* that keeps local copies of things recently accessed from the shared memory. The use of individual cache memories at each processor radically reduces traffic to the shared memory.

The shared memory may be a single physical memory unit, as in the SUN SPARCCenter. This kind of computer organization cannot be scaled indefinitely upward—the shared memory and its bus become a bottleneck at some point.

A more scalable *distributed memory* design has a single shared memory address space, but the physical memory is distributed among the nodes. This arrangement exploits microprocessors' low memory cost and gives better performance for local access. Many experts believe this will become the dominant organization for machines with more than a few processors.

Some distributed-memory machines, such as the Convex Exemplar, enforce cache coherence, so that all processors see the same memory values. Others, such as the Cray T3D, do not enforce coherence but use very fast special circuits to get very low shared-memory latency. Most machines with a shared physical memory maintain cache coherence.

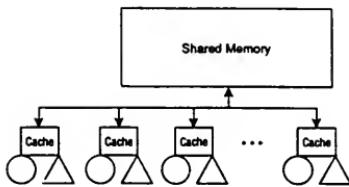


FIGURE A.4 A shared-memory parallel computer organization.

Generations of Parallel Computers

First Commercial Generation: SIMD

Parallel computers with small numbers of processors have been standard commercial fare for 30 years. In some cases, the multiple processors were in duplex, triplex, or quadruplex configurations for high availability; in most advanced computers there have been processors dedicated to input-output. Most vector computers have also been modestly parallel for more than a decade. One-of-a-kind highly parallel computers have been built now and then since the 1960s, with limited success. The Advanced Research Projects Agency (ARPA) recognized the technical-economic imperative to develop highly parallel computers for both military and civilian applications and acted boldly to create its high-performance computing program. This stimulus combined with a ferment of new ideas and with entrepreneurial enthusiasm to encourage several manufacturers to market highly parallel machines, among them Intel, Ncube, Thinking Machines Corporation (TMC), and MasPar. Most of these first-generation machines were SIMD computers, exemplified by the CM1 (Connection Machine 1) developed by Thinking Machines.

Because SIMD execution lacks the information content of multiple instruction flows, applications have to be more uniform to run efficiently on SIMD computers than on other types of parallel computers. Compounding this inherent difficulty, the first-generation machines had only primitive software tools. No application software was available off the shelf, and existing codes could not be automatically ported, so that each application had to be rebuilt from start. Moreover,

few of the first-generation machines used off-the-shelf microprocessors with their economic and software advantages.

The first generation of highly parallel computers had some successes but proved to be of too limited applicability to succeed in the general market. Some naturally parallel applications were reprogrammed for these machines, realizing gains in execution speed nearly in proportion to the number of processors applied to the problem, up to tens of processors. The set of applications for which this was true was quite limited, however, and most experts agree that the SIMD configuration has its units too tightly coupled to be used effectively in a wide variety of applications. Nonetheless, the creation of this generation of machines, and their provision of a platform for pioneering and experimental applications, clearly started a great deal of new thinking in academia about how to use such machines.

Second Generation: Message-Passing MIMD

Striving for the wider applicability that would be enabled by a more flexible programming style, parallel computer researchers and vendors developed MIMD configurations made up of complete microprocessors (sometimes augmented by SIMD clusters). By and large, these machines used message-passing for interprocessor communication. The Thinking Machines CM5 is a good example of this second generation. Other examples use off-the-shelf microprocessors as nodes.

Although improving somewhat in ease of use, such machines are still hard to program, and users still need to change radically how they think and the type of algorithms they use. Moreover, because these machines were different both from conventional computers and from first-generation highly parallel computers, the compilers and operating systems again had to be redone "from scratch" and were primitive when the machines were delivered.

The second-generation machines have proven to be much more widely applicable, but primitive operating systems, the continuing lack of off-the-shelf applications, and the difficulties of programming with elementary tools prevented widespread adoption by computer-using industries. As the market registered its displeasure with these inadequacies, several of the vendors of first- and second-generation parallel computers, including TMC and Kendall Square Research, went into Chapter 11 protection or retired from the parallel computer-building field. A beneficial side effect of these collapses has been the scattering of parallel-processing talent to other vendors and users.

As parallel computers gained acceptance, existing vector computer vendors claimed their sales were being harmed by the government promotion and subsidization of a technology that they saw as not yet ready to perform. Cray Research and Convex, among others, saw their sales fall, partly due to performance/cost breakthroughs in smaller computers, partly due to the defense scale-back, and partly due to some customers switching from vector to parallel computers. The complaints of the vector computer vendors triggered studies of the HPCCI by the General Accounting Office and the Congressional Budget Office (see "Concerns Raised in Recent Studies"). Cray Research and Convex have since become important vendors of parallel computers.

Third Generation: Memory-Sharing MIMD

In the third generation, major existing computer manufacturers independently decided that the shared-memory organization, although limited in ultimate scalability, offered the most viable way to meet present market needs. Among others, SGI, Cray Research, and Convex have made such systems using off-the-shelf microprocessors from MIPS, IBM, DEC, and Hewlett-Packard, respectively. As noted above, market acceptance has been encouraging—industrial computer users have been buying the machines and putting them to work. Many users start by using standard

software and running the systems as uniprocessors on bread-and-butter jobs, and then expand the utilization of the multiple processors gradually. As parallel algorithms, compilers, languages, and tools continue to develop, these memory-shared machines are well positioned to capitalize on them.

Programming

The development of parallel computing represents a fundamental change not only in the machines themselves, but also in the way they are programmed and used. To use fully the power of a parallel machine, a program must give the machine many independent operations to do simultaneously, and it must organize the communication among the processor nodes. Developing techniques for writing such programs is difficult and is now regarded by the committee as the central challenge of parallel computing.

Computer and computational scientists are now developing new theoretical concepts and underpinnings, new programming languages, new algorithms, and new insights into the application of parallel computing. While much has been done, much remains to be done: even after knowledge about parallel programming is better developed, many existing programs will need to be rewritten for the new systems.

Algorithms

There is a commonly held belief that our ability to solve ever larger and more complex problems is due primarily to hardware improvements. However, A.G. Fraser of AT&T Bell Laboratories has observed that for many important problems the contributions to speed-ups made by algorithmic improvements exceed even the dramatic improvements due to hardware. As a long-term example, Fraser cited the solution of Poisson's equation in three dimensions on a 50 by 50 by 50 grid. This problem, which would have taken years to solve in 1950, will soon be solved in a millisecond. Fraser has pointed out that this speed-up is owing to improvements in both hardware and algorithms, with algorithms dominating.¹

During the mid-1980s, several scientists independently developed tree codes or hierarchical N-body codes to solve the equations of the gravitational forces for large multibody systems. For 1 million bodies, tree codes are typically 1,000 times faster than classic direct-sum algorithms. More recently, some of these tree-code algorithms have been modified to run on highly parallel computers. For example, Salmon and Warren have achieved a speed-up of 445 times when running their codes on a computer with 512 processors as compared with running them on a single processor (Kaufman and Smart, 1993, pp. 73-74).

Over the half-century that modern computers have been available, vast improvements in problem solving have been achieved because of new algorithms and new computational models; a short list from among the numerous examples includes:

- Finite-element methods,
- Fast Fourier transforms,
- Monte Carlo simulations,
- Multigrid methods,
- Methods for sparse problems,
- Randomized algorithms,
- Deterministic sampling strategies, and
- Average case analysis.

The exponential increase in the sizes of economical main memories has also enabled a host of new table-driven algorithms, techniques that were unthinkable a decade ago. Discovering and developing new algorithms for solving both generic and specific problems from science, engineering, and the financial services industry, designed and implemented on parallel architectures, will continue to be an important area for national investment.

A SKETCH OF THE HPCCI'S HISTORY

Development and Participants

To quote from the 1993 Blue Book: "The goal of the federal High Performance Computing and Communications Initiative (HPCCI) is to accelerate the development of future generations of high-performance computers and networks and the use of these resources in the federal government and throughout the American economy" (FCCSET, 1992). This goal has grown, like the HPCCI itself, from many roots and has continued to evolve as the initiative has matured. Box A.2 illustrates the evolution of the HPCCI's goals as presented by the Blue Book annual reports.

BOX A.2 HPCCI Goals As Stated in the Blue Books

FY 1992

Accelerate significantly the commercial availability and utilization of the next generation of high-performance computers and networks:

- Extend U.S. technological leadership in high-performance computing and communications.
- Widely disseminate and apply technologies to speed innovation and to serve the national economy, national security, education, and the global environment.
- Spur productivity and competitiveness.

FY 1993

Unchanged.

FY 1994

Goals remained the same with addition of the Information Infrastructure and Technology Applications and program element and mention of the National Information Infrastructure.

FY 1995

Meta-goal ("Accelerate significantly . . . ") not mentioned. Goals consolidated as:

- Extend U.S. technological leadership in high-performance computing and communications; and
- Widely disseminate and apply technologies to speed innovation and to improve national economic competitiveness, national security, education, health care (medicine), and the global environment.

Beginning in the early 1980s, several federal agencies advanced independent programs in high-performance computing and networking.² The National Science Foundation (NSF) built on recommendations from the National Science Board Lax report in 1982,³ as well as a set of internal reports⁴ that recommended dramatic action to end the 15-year supercomputer famine in U.S.

universities. NSF asked Congress in 1984 for funds to set up, by a national competition, a number of supercomputer centers to provide academic researchers access to state-of-the-art supercomputers, training, and consulting services. Very quickly this led to the creation of an NSF network backbone to connect the centers. This in turn provided a high-speed backbone for the Internet. Several organizations, including the Office of Management and Budget and the former Federal Coordinating Council on Science, Engineering, and Technology (FCCSET) of the Office of Science and Technology Policy (OSTP), built on these activities and similar efforts in the Department of Energy (DOE),⁵ the National Aeronautics and Space Administration (NASA), and the Department of Defense (DOD) to develop the concept of a National Research and Education Network (NREN) program (CSTB, 1988). These explorations were linked to other concurrent efforts to support advanced scientific computing among researchers and to promote related computer and computational science talent development. The result was the High-Performance Computing Program. The program included an emphasis on communications technology development and use from the outset.

High-performance Computing Program structure and strategy were discussed intensively within several federal agencies in 1987-1988, resulting in initial formalization and publication of a program plan in 1989 (OSTP, 1989). OSTP provided a vehicle for interagency coordination of high-performance computing and communications activities, acting through FCCSET and specific subgroups, including the Committee on Physical, Mathematical, and Engineering Sciences; its subordinate Panel on Supercomputers; its High Performance Committee (later subcommittee); its Research Committee (later subcommittee); and its High Performance Computing, Communications, and Information Technology (HPCCIT) Subcommittee. The initial HPCCI effort was concentrated in four agencies: DOD's Advanced Research Projects Agency, DOE, NASA, and NSF. These agencies remain the dominant supporters of computing and computational science research. Although not then a formal member, the National Security Agency (NSA) has also always been an influential player in high-performance computing, due to its cryptography mission needs.

High-performance computing activities received added impetus and more formal status when Congress passed the High-Performance Computing Act of 1991 (PL 102-194) authorizing a 5-year program in high-performance computing and communications. This legislation affirmed the interagency character of the HPCCI, assigning broad research and development (R&D) emphases to the 10 federal agencies that were then participating in the program without precluding the future participation of other agencies. The group of involved agencies expanded to include the Environmental Protection Agency, National Library of Medicine (part of the National Institutes of Health), National Institute of Standards and Technology (part of the Department of Commerce (DOC), and National Oceanographic and Atmospheric Administration (part of DOC) as described in the 1992 and 1993 Blue Books. Additional agencies involved subsequently include the Education Department, NSA, Veterans Administration (now the Department of Veteran Affairs), and Agency for Health Care Policy and Research (part of the Department of Health and Human Services). These and other agencies have participated in HPCCIT meetings and selected projects either as direct members or as observers.

Since its legislative inception in 1991, the HPCCI has attained considerable visibility both within the computer research community and as an important element of the federal government's technology programs. When originally formulated, the HPCCI was aimed at meeting several "Grand Challenges" such as modeling and forecasting severe weather events. It was subsequently broadened to address "National Challenges" relating to several important sectors of the economy, such as manufacturing and health care, and then the improvement of the nation's information infrastructure. The evolution of emphasis on the Grand and National Challenges and also the nation's information infrastructure is outlined in Box A.3.

BOX A.3 From Grand Challenges to the National Information Infrastructure and National Challenges: Evolution of Emphasis as Documented in the Blue Books

FY 1992

- Grand Challenges featured in title and discussed in text
 - Forecasting severe weather events
 - Cancer gene research
 - Predicting new superconductors
 - Simulating and visualizing air pollution
 - Aerospace vehicle design
 - Energy conservation and turbulent combustion
 - Microelectronics design and packaging
 - Earth biosphere research
- National Challenges not discussed

FY 1993

- Grand Challenges featured in title and discussed in text
 - Magnetic recording technology
 - Rational drug design
 - High-speed civil transports (aircraft)
 - Catalysis
 - Fuel combustion
 - Ocean modeling
 - Ozone depletion
 - Digital anatomy
 - Air pollution
 - Design of protein structures
 - Venus imaging
 - Technology links to education
- National Challenges not discussed

FY 1994

- National Information Infrastructure (NII) featured in title
 - Medical emergency and weather emergency discussed as examples of potential use of NII
- Potential National Challenge areas listed in Information Infrastructure Technology and Applications discussion
 - Civil infrastructure
 - Digital libraries
 - Education and lifelong learning
 - Energy management
 - Environment
 - Health care
 - Manufacturing processes and products
 - National security
 - Public access to government information

continues

BOX A.3—continued

- Grand Challenges discussed as case studies in text
 - Climate modeling
 - Sharing remote instruments
 - Design and simulation of aerospace vehicles
 - High-performance life science: molecules to magnetic resonance imaging
 - Nonrenewable energy resource recovery
 - Groundwater remediation
 - Improving environmental decision making
 - Galaxy formation
 - Chaos research and applications
 - Virtual reality technology
 - High-performance computing and communications and education
 - Guide to available mathematics software
 - Process simulation and modeling
 - Semiconductor manufacturing for the 21st century
 - Field programmable gate arrays
 - High-performance Fortran and its environment

FY 1995

- National Information Infrastructure featured in title and discussed in text
 - Information infrastructure services
 - Systems development and support environments
 - Intelligent interfaces
- National Challenge areas discussed in text
 - Digital libraries
 - Crisis and emergency management
 - Education and lifelong learning
 - Electronic commerce
 - Energy management
 - Environmental monitoring and waste minimization
 - Health care
 - Manufacturing processes and products
 - Public access to government information
- Major section devoted to "High-Performance Living" with future scenario based on the National Challenges and the National Information Infrastructure
- Grand Challenges discussed in text. More than 30 Grand Challenges illustrated by examples within the following larger areas:
 - Aircraft design
 - Computer science
 - Energy
 - Environmental monitoring and prediction
 - Molecular biology and biomedical imaging
 - Product design and process optimization
 - Space science

Concerns Raised in Recent Studies

A 1993 General Accounting Office (GAO) study of ARPA activities related to the HPCCI and a 1993 Congressional Budget Office (CBO) study of HPCCI efforts in massively parallel computing have been regarded by some as being critical of the entire HPCCI. The committee, which received detailed briefings from the studies' authors, offers the following observations.

GAO Report

The GAO report⁶ did not attempt to evaluate the entire HPCCI but focused instead on research funding, computer prototype acquisition activities, and the balance between hardware and software investments by ARPA. It recommended that ARPA (1) broaden its computer placement program by including a wider range of computer types, (2) establish and maintain a public database covering the agency HPCCI program and the performance characteristics of the machines it funds, and (3) emphasize and provide increased support for high-performance software. The report's authors stated to the committee that although recommending improvements, they had found that ARPA had administered its program with propriety.

The committee notes that progress has been made on each of GAO's recommendations, and it has urged that further progress be supported. Committee recommendation 4 calls for further reduction in funding of computer development by vendors and for experimental placement of new machines. These actions should result in a wider variety of machine types as agencies select different machines to meet their mission needs. The National Coordination Office (NCO) has made more program information available and the committee recommends that functions in this area receive added attention by an augmented NCO (recommendation 11). Likewise the committee has called in recommendation 3 for added emphasis on the development of software and algorithms for high-performance computing.

CBO Report

The primary theme of the CBO report (1993) was that because it was aimed primarily at massively parallel machines, which currently occupy only a small part of the computer industry, the High-Performance Computing Systems component of the HPCCI would have little impact on the computer industry. (The high-performance communications and networking segment of the program is not addressed in the CBO report.) The CBO report assumed that the HPCCI was to support the U.S. computer industry, in particular the parallel-computing portion. Although this might be an unstated objective, the explicitly stated goals relate instead to developing new high-performance computer architectures, technologies, and software. The HPCCI appears to be fulfilling the stated goals.

The CBO report did not attempt to analyze the impact of the development of high-performance computing and communications technology on the larger computer industry over a longer period of time. The primary focus of the high-performance computing portion of the program is the creation of scalable parallel machines and software. It is widely believed in both the research and industrial communities that parallelism is a key technology for providing long-term growth in computing performance, as discussed in the early sections of this appendix. The HPCCI has demonstrated a number of successes in academia, in industry, and in government laboratories that provide a significant increase in our ability to build and use parallel machines. Just as reduced instruction set computers (RISC) technology, developed partly with ARPA funding, eventually became a dominant force in computing (some 10 years after the program started), the initiative's

ideas are starting to take root in a larger context across the computer industry. Since using parallel processors requires more extensive software changes than did embracing RISC concepts, we should expect that multiprocessor technology will take longer to be adopted.

NOTES

1. A.G. Fraser, AT&T Bell Laboratories, personal communication.
2. Department of Energy (DOE) officials point out that their efforts date from the mid-1970s. For example, in 1974 DOE established a nationwide program providing energy researchers with access to supercomputers and involving a high-performance communications network linking national laboratories, universities, and industrial sites, the precursor of today's Energy Sciences Network (ESNet). See Nelson (1994).
3. *Report of the Panel on Large Scale Computing in Science and Engineering*, Peter Lax, Chairman, sponsored by the U.S. Department of Defense and the National Science Foundation, in cooperation with the Department of Energy and the National Aeronautics and Space Administration, Washington, D.C., December 26, 1982.
4. *A National Computing Environment for Academic Research*. Marcel Bardon and Kent Curtis, NSF Working Group on Computers for Research. National Science Foundation, July 1983.
5. The DOE laboratories had been involved in supercomputing since World War II and were not particularly affected by the setting up of the NSF centers or FCCSET until the late 1980s or early 1990s as the HPCCI emerged.
6. See GAO (1993). Another GAO report (1994) was not released in time for the committee to receive a detailed briefing on which to base further group deliberations. However, observations from that report are drawn in the body of this report.

B

High-Performance Communications Technology and Infrastructure

HIGH-PERFORMANCE COMMUNICATIONS TECHNOLOGY AND INFRASTRUCTURE ADVANCE

The performance/cost imperatives of communications have driven the technology from parallel to serial, from hundreds of slower wires to a few very fast fibers. Fiber-optic transmission offers stunning bandwidths: 100,000 telephone calls or 800 video channels on one pair of fibers. Communications is different from computing. Value often comes from whom one can talk to, rather than how rapidly. Issues of scaling are very important. The scale of the needed networking raises a host of new research issues as to how millions of users can attach to the network.

High-performance computing and high-performance communications support each other in complex ways. Communications has become digital, and the switching of fast digital signals requires high-performance computing technology. On the other hand, very fast computer-to-computer communications are crucial for many applications. Today 16 percent of investment in the High Performance Computing and Communications Initiative (HPCCI) is directed at communications. The communications content of the HPCCI has two aspects: research and development to advance communications and related capabilities (see Table B.1), and delivery of access to communications-based infrastructure to researchers to facilitate their work (see Table B.2). The introduction of the fifth HPCCI component, Information Infrastructure Technology and Applications (IITA), in 1993 appears to extend the second aspect: awards and activities associated with this component appear to emphasize making existing capabilities more useful and more widely used, as opposed to developing new communications-based capabilities, which appears to be largely supported under the National Research and Education Network (NREN) component of the HPCCI.

The Internet is the centerpiece of the present HPCCI communications infrastructure program; it includes the network elements (backbone, regional, and "connections") supported under the NREN program. Thanks to federal support of internetworking technologies that have been applied in the Internet generally and specifically in NREN-supported elements (NSFNET, ESnet, and the NASA Science Internet), the United States has a strong lead in these technologies worldwide. The United States is home to a vital industry that supplies related equipment and software, including businesses begun as spinoffs from academic research activity. See Box B.1 for an example of how government, academic, and industrial investments can complement each other to accelerate the development of a key communications technology.

TABLE B.1 HPCCI Program Activities in Communications Research, FY 1995

Component ^a	Agency ^b	Funding Request for FY 1995 (millions of dollars)	Activity
NREN	ARPA	43.10	Networking
IITA	ARPA	23.00	Global grid communications
BRHR	NSF	11.30	Very high speed networks and optical systems
NREN	NSA	03.50	Very high speed networking
NREN	NSA	02.60	High-speed data protection electronics
NREN	DOE	02.00	Gigabit research and development
NREN	NIST	01.75	Metrology to support mobile and fixed-base communications networks

^aNREN, National Research and Education Network; IITA, Information Infrastructure Technology and Applications; BRHR, Basic Research and Human Resources.

^bARPA, Advanced Research Projects Agency; NSF, National Science Foundation; NSA, National Security Agency; DOE, Department of Energy; NIST, National Institute of Standards and Technology.

TABLE B.2 HPCCI Program Activities in Communications Infrastructure, FY 1995

Component ^a	Agency ^b	Funding Request for FY 1995 (millions of dollars)	Activity ^c
NREN	NSF	46.16	NSFNET
NREN	DOE	14.80	Energy sciences network (ESnet)
NREN		12.70	NREN
NREN	NOAA	08.70	Networking connectivity
NREN	NIH	06.50	NLM medical connections program
NREN	NIST	02.20	NREN deployment and performance measures for gigabit nets and massively parallel processor systems
IITA	NIH	02.00	NCI high-speed networking and distributed conferencing
NREN	EPA	00.70	State network connectivity

^aNREN, National Research and Education Network; IITA, Information Infrastructure Technology and Applications.

^bNSF, National Science Foundation; DOE, Department of Energy; NOAA, National Oceanic and Atmospheric Administration; NIH, National Institutes of Health; NIST, National Institute of Standards and Technology; EPA, Environmental Protection Agency

^cNLM, National Library of Medicine; NCI, National Cancer Institute.

**BOX B.1 Federal Government Participation
in the Development of Asynchronous Transfer Mode**

The federal government, through the HPCCI and other programs, played a part in the development of the asynchronous transfer mode, or ATM, standard, and most importantly, played a very significant role in developing broad support for ATM as an important switching technology for high-speed computer networks. ATM was developed as a switching technology by the telecommunications community for application in the so-called broadband integrated services digital network, or BISDN. However, its development did not occur without some controversy.

Telecommunications switching had always been based on circuit-switching technologies, which are very well suited to providing fixed-rate connections in a network. Multirate circuit switching can be used to provide connections that are any multiple of some basic rate. Narrowband ISDN is based on these technologies, and most switching and transmission experts in the telecommunications industry expected a straightforward extension of ISDN circuit-switching technology, based on multiples of a 64-Kbps basic rate, into the broadband realm of speeds in the 155-Mbps range.

Variable rate communications services were clearly needed for data communications, as demonstrated by the ARPANET in the 1970s. However, it was also recognized that video and even voice services might be provided in a more efficient manner than was possible through circuit switching if the basic network service could support variable rate communications. Between 1984 and 1987 researchers in leading telecommunications and academic laboratories developed a number of fast-packet switching, or FPS, systems that laid the technological groundwork for the ATM standard. A noted academic researcher participating in this effort was Jonathan Turner, whose work was supported by both NSF and industry funding. Telecommunications industry laboratories such as AT&T, Bellcore, CNET (France), and Bell Telephone Manufacturing (Belgium) also played a significant role. Bellcore, for example, built the first prototype broadband central office in late 1986, and it contained a packet switch that operated at 50 Mbps per line. By 1987, this rate had been extended to nearly 155 Mbps per line. These efforts, in essence, proved that packet switching was capable of running at the data rates required by BISDN. At this time, ARPANET and NSFNET were running at speeds of only 64 Kbps to 1.5 Mbps, and local area networks such as Ethernet ran at speeds of only 10 Mbps.

In late 1986, Gordon Bell, then at NSF, visited Bellcore and received a briefing on fast CMOS-based packet switching technology that had been simulated at speeds of 150 Mbps. Several months later he called a workshop of computer networking researchers to discuss the future of high-speed computer networking. Out of that workshop, from David Farber (University of Pennsylvania) and Bob Kahn (Corporation for National Research Initiatives), came a mid-1987 proposal to NSF to form the gigabit testbed projects.

The telecommunications industry was quite active during this time in developing a standard based on packet switching. In 1988, the Consultative Committee on International Telephony and Telegraphy selected (and named) ATM as the standard for BISDN switching. This action served to focus the telecommunications industry on the development of this technology. However, in 1988 ATM was little more than a name, a basic packet format, and a common cause across the telecommunications industry alone.

In 1989, the gigabit testbed projects were formally initiated with IDARPA and NSF funding. One of the objectives of the gigabit testbeds was to determine what switching technologies were appropriate for use in gigabit networking. There were several major contenders: ATM, HiPPi (a switching technology favored by the supercomputer community), and PTM (a proprietary packet-switching technology advocated by IBM, different from ATM primarily in that it used variable-sized packets). Although it would be several years before conclusive results from the testbeds were produced, the use of ATM in the testbeds gave it wide exposure within the federal government's technology community.

continues

BOX B.1—continued

In 1990, the first company devoted solely to producing ATM products was formed. The founders of FORE Systems were first exposed to ATM technology through their participation in the NECTAR gigabit testbed. Cross-testbed exchanges were encouraged by the gigabit testbed program, and in 1989 and 1990 there were meetings of the Bellcore and MIT researchers working on ATM for Aurora and the Carnegie Mellon University researchers working on NECTAR. FORE was founded in 1990 and subsisted through its first product development cycle largely on research and prototyping contracts from ARPA and the Naval Research Laboratory. FORE's first products were delivered to the Naval Research Laboratory in 1991. FORE today is a rapidly growing company with nearly 300 employees and one of the leading manufacturers of ATM switching equipment.

In 1990, a collaboration was formed between Apple, Bellcore, Sun Microsystems, and Xerox to attempt to gain the acceptance of ATM as a new local area networking standard. This group, which eventually expanded to include Digital and Hewlett-Packard, produced the first published specification of ATM for Local Networking Applications in April 1991. ARPA participated in informal discussions with this group via its ARPA Networking Principal Investigator meetings, looking for opportunities to fund technical work that would be necessary to make ATM useful for local area networking. One notable result of ARPA's funding was the first implementation of the Q.2931 ATM signaling standard, which was required to allow ATM to implement switched (as opposed to permanent) virtual circuits. Another result was the implementation, by Xerox, of an ATM switch architecture that is well suited for low-cost implementation. This work, funded in part by ARPA, served to accelerate the development of an industry standard for local ATM by the ATM Forum.

Once ATM had been established as both a computer and telecommunications networking standard, the role of the federal government turned toward developing an early market for high-speed networking equipment and services. Several notable programs greatly expanded this early market. One has been the gigabit testbeds. Some of the first installations of 2.4-Gbps SONET and ATM by several telecommunications carriers were for these testbeds. Although the federal government did not spend a single dollar for these facilities, its leadership of the gigabit testbed program was critical in getting telecommunications carriers to construct these very expensive facilities for the testbeds. Telecommunications equipment manufacturers benefited directly and saw their development of very high speed (full-duplex 622-Mbps) ATM equipment greatly accelerated.

Some direct federal procurements of ATM and SONET services are still playing a key role in the development of the marketplace. An important procurement of ATM and SONET services and technology was the Washington Area Bitway (WABitway), the first significant sale of high-speed SONET services by Bell Atlantic. In early 1994, NSF announced the selection of a number of companies to provide ATM-based communications services for the new NSFNET. Telecommunications companies involved in this project include Ameritech, MCI, MFS Datanet, Pacific Bell, and Sprint.

Eight HPCCI program activities are directed primarily at communications infrastructure, principally supporting deployment of the Internet within each agency's community. Several of the infrastructure programs are building on early results of the gigabit testbed research. For example, ATM and SONET networking technologies, first deployed in the gigabit testbeds in 1992, appear in some form in many of the FY 1995 infrastructure activities.

Developing and broadening access to information infrastructure pose many research issues. Information infrastructure is more complex than networks, per se, and the computing and communications research community has already helped to explore and define fundamental concepts, for example, the concept of "middleware" to cover the kind of internal services that help to transform a network into information infrastructure. Research into how to implement such services has begun under the HPCCI umbrella. More specifically, the National Aeronautics and Space Administration, the National Science Foundation, and the Advanced Research Projects Agency have combined to fund research to support the development of digital libraries, providing a vehicle for exploring many concepts associated with information infrastructure (NSF, 1993).

The concept of a better nationwide information infrastructure, itself connected to a global information infrastructure, poses yet other concerns associated with interconnecting multiple kinds of networks from multiple kinds of providers to multiple kinds of users offering multiple kinds of services. This construct adds great complexity, increasing the emphasis on scale and architecture and adding in such concerns as heterogeneity of systems, decentralization of control and management, routing, security, and so on. There have been many government, academic, and industry studies under way to identify and clarify these research issues. Although the newspapers are filled with announcements of corporate alliances, new venture formations, and new product introductions more or less linked to the advancement of the nation's information infrastructure, significant advances call for the solution of many technical problems and therefore for a significant research effort.

C Review of the High Performance Computing and Communications Initiative Budget

The committee attempted a nontraditional look at how High Performance Computing and Communications Initiative (HPCCI) funds are being invested.¹ Traditional HPCCI budget reports show budget breakdowns by agency and by program component (High-Performance Computing Systems, National Research and Education Network, Advanced Software Technology and Algorithms, Information Infrastructure Technology and Applications, and Basic Research and Human Resources). The committee found it informative to examine the funding from a functional perspective to understand what sort of technical work is being performed and in what quantity.

The committee separated the 88 HPCCI program elements into 11 disciplines defined as indicated below:

- *Computer technology (CPT)*—applied research directed at advancing the state of computer architecture and hardware technology;
- *Software technology (SWT)*—applied research directed at advancing the state of computer software technology;
- *Communications technology (CMT)*—applied research directed at advancing the state of communications technology;
- *Computing infrastructure (CPI)*—acquisition and operation of supercomputer facilities;
- *Communications infrastructure (CMI)*—acquisition and operation of high-performance computer communications networks and services;
- *Applications and computational science (APP)*—creation of software and computational techniques directed at solving specific scientific problems and applications;
- *Common applications support (CAS)*—creation of software and computational techniques to support a range of applications across multiple disciplines;
- *Artificial intelligence and human-machine interaction (AI)*—applied research directed at solving artificial intelligence and human interface problems.

- *Basic hardware technology (BHW)*—basic electronics research supporting electronic components that might be applied to a wide variety of systems, including computers and communications systems;
- *Education (EDU)*—training and education; and
- *Administration (ADM)*—National Coordination Office.

The committee classified 86 of the 88 program elements as coming under 1 of the 11 disciplines listed above, based on each program element's FY 1995 milestones (Table C.1). If a program element appeared to fit into more than one discipline, the committee categorized it by examining the element's milestones to determine where the majority of the program activity was concentrated. Two of the larger Advanced Research Projects Agency (ARPA) program activities (Intelligent Systems and Software, and Information Sciences) were split between two disciplines.

Table C.2 shows the FY 1993 actual budget, the FY 1995 request,² and the percentage change in the HPCCI budget for each of these 11 disciplines.

BUDGET REVIEW

It is interesting to examine the HPCCI budget to see which areas are being emphasized and to compare these with the HPCCI's goals and objectives. As indicated also in Chapter 2, the current program goals are as follows:

- Extend U.S. leadership in high-performance computing and networking technologies;
- Disseminate the technologies to accelerate innovation and serve the economy, national security, education, and the environment; and
- Spur gains in U.S. productivity and industrial competitiveness.

The computer technology, software technology, and communications technology disciplines address the goals of extending technical leadership in computing and communications and providing key enabling technologies for the information infrastructure. The budget for these three disciplines accounted for 32.9 percent of the FY 1993 actual budget and 30.5 percent of the FY 1995 requested budget.

The largest part of the HPCCI budget is invested in applications and supercomputer computing infrastructure to support applications—49.8 percent of the FY 1993 actual budget and 50.1 percent of the FY 1995 requested budget. Applications and computational science, common applications support, artificial intelligence and human-machine interaction, and computing infrastructure programs are included. The rest of the budget requested for FY 1995 is divided among basic hardware technology (5.1 percent), communications infrastructure (7.9 percent), education (5.6 percent), and a very small amount for program administration.

TABLE C.1 Mapping of Agencies' HPCCI Budget Applications, Program Elements, and Discipline Categories

Agency ¹	1993 Requested (\$M)	1993 Actual (\$M)	Discipline ¹	Program Element
ARPA	36.15	11.35	AI	Intelligent systems and software
ARPA	36.15	11.35	SWT	
ARPA	60.20	44.90	CPT	Scalable computing systems
ARPA	44.50	33.50	BHW	Microsystems
ARPA	43.10	34.80	CMT	Networking
ARPA	33.90	38.00	SWT	National-scale information enterprises
ARPA	29.60	36.50	SWT	Scalable software
ARPA	23.00	00.00	CMT	Global grid communications
ARPA	10.50	15.10	AI	Information sciences
ARPA	10.50	15.10	SWT	
ARPA	14.00	13.90	EDU	Foundations
ARPA	09.80	00.00	APP	Health information infrastructure
ARPA	06.00	00.00	BHW	Integrated command and control technology
NSF	76.43	63.89	CPI	Supercomputer centers
NSF	46.16	30.10	CMI	NSFNET
NSF	35.25	00.00	CAS	Information infrastructure technology and applications program
NSF	25.35	21.79	SWT	Software systems and algorithms
NSF	20.95	18.65	CPI	Research infrastructure
NSF	20.70	18.76	CPT	Computing systems and components
NSF	20.24	15.34	EDU	Education and training
NSF	11.50	07.80	APP	Biological sciences (non-NC/GC)
NSF	11.30	09.80	CMT	Very high speed networks and optical systems
NSF	11.00	10.40	AI	Human machine interaction and information access
NSF	10.75	07.00	CAS	Grand Challenge applications groups
NSF	10.55	09.20	CAS	Research centers
NSF	09.77	02.75	APP	Physical sciences (non-NC/GC)
NSF	07.59	05.72	CAS	Computational mathematics (non-NC/GC)

continues

TABLE C.1—*continued*

Agency ^a	1995 Requested (\$M)	1993 Actual (\$M)	Discipline ^b	Program Elements ^c
NSF	04.23	02.17	APP	Engineering (non-NC/GC)
NSF	03.84	01.15	CPI	Geosciences (non-NC/GC)
NSF	03.01	00.65	APP	Social, behavioral, and economic sciences (non-NC/GC)
DOE	35.60	35.13	CPI	Supercomputer access
DOE	16.00	15.25	CAS	Basic research for applied mathematics research
DOE	14.80	07.68	CMI	Energy sciences network (ESnet)
DOE	12.90	07.15	CPI	High-performance computing research centers
DOE	12.60	08.98	CAS	Software components and tools
DOE	09.90	07.73	CPI	Evaluation of early systems
DOE	09.00	06.53	APP	Enabling energy Grand Challenges
DOE	03.40	02.44	CAS	Computational techniques
DOE	03.00	02.36	EDU	Education, training, and curriculum
DOE	02.00	02.60	EDU	Research participation and training
DOE	02.00	01.86	CMT	Gigabit research and development
DOE	02.00	01.80	APP	High-performance research centers—global climate collaboration
DOE	01.20	00.00	CAS	Information infraservices
DOE	01.00	00.00	CPI	Advanced prototype systems
NASA	55.30	46.80	APP	Grand Challenge support
NASA	26.40	17.60	CPI	Testbeds
NASA	12.70	08.50	CMI	National Research and Education Network (NREN)
NASA	10.70	00.00	EDU	Information infrastructure applications
NASA	09.20	05.40	SWT	Systems software
NASA	06.80	00.00	CAS	Information infrastructure technology
NASA	03.80	03.30	EDU	Basic research and human resources
NIH	11.00	08.00	CPI	DCRT high-performance biomedical computing program

continues

TABLE C.1—*continued*

Agency ^a	1995 Request (\$M)	1993 Actual (\$M)	Discipline ^b	Program ^c
NIH	11.00	01.50	APP	National Library of Medicine high-performance computing and communications health care applications
NIH	08.80	03.40	APP	NCRR information infrastructure technology applications
NIH	08.80	06.80	APP	NCRR advanced software technology and algorithms
NIH	06.70	06.30	CPI	NCI Frederick biomedical supercomputing center
NIH	06.50	00.40	CMI	NLM medical connections program
NIH	05.40	03.80	CAS	NLM IAIMS grants
NIH	05.00	03.10	EDU	NCRR basic research and human resources
NIH	04.80	04.10	APP	NLM biotechnology informatics
NIH	04.70	05.00	CAS	NLM intelligent agent database searching
NIH	03.60	02.90	EDU	NLM HPCCI training grants
NIH	02.20	01.50	APP	NLM electronic imaging
NIH	02.00	00.00	CMI	NCI high-speed networking and distributed conferencing
NIH	00.70	00.40	ADM	National Coordination Office
NIH	00.60	00.00	APP	High-performance communications for PDQ, Cancer Net, and electronic publishing
NSA	26.10	00.00	CPT	Supercomputing research
NSA	05.70	00.00	SWT	Secure operating system development
NSA	03.50	00.00	CMT	Very high speed networking
NSA	02.60	00.00	CMT	High-speed data protection electronics
NSA	02.00	00.00	BHW	Superconducting research
NSA	00.23	00.00	EDU	Technology-based training
NIST	25.20	00.00	CAS	Systems integral for manufacturing applications
NIST	07.60	00.60	APP	Development and dissemination of scientific software for high-performance computing systems

continues

TABLE C.1—*continued*

Agency ^a	1995 Requested (\$M)	1993 Actual (\$M)	Discipline ^b	Program Element ^c
NIST	06.45	00.00	BHW	Metrology for future generations of microelectronics
NIST	04.00	00.00	CAS	Language, image, and text processing
NIST	04.00	00.00	SWT	Specification and testing of high-integrity, distributed systems
NIST	02.75	00.00	CAS	Support for electronic commerce
NIST	02.20	01.50	CMT	Deployment and performance measures for gigabit nets and massively parallel processor systems
NIST	01.75	00.00	CMT	Metrology to support mobile and fixed-base communications networks
NIST	01.25	00.00	CAS	Electronic libraries and distributed multimedia applications
NIST	01.20	00.00	SWT	Assurance, reliability, and integrity of NREN objects
NOAA	16.05	09.40	APP	Advanced computation
NOAA	08.70	00.40	CMI	Networking connectivity
NOAA	00.50	00.00	APP	Information dissemination pilots
EPA	06.45	05.33	APP	Environmental modeling
EPA	05.25	01.31	APP	Computational techniques
EPA	01.97	01.16	EDU	Education and training
EPA	00.70	00.21	CMI	State network connectivity
EPA	00.30	00.00	APP	Public data access

^aARPA, Advanced Research Projects Agency; NSF, National Science Foundation; DOE, Department of Energy; NASA, National Aeronautics and Space Administration; NIH, National Institutes of Health; NSA, National Security Agency; NIST, National Institute of Standards and Technology; NOAA, National Oceanographic and Atmospheric Administration; EPA, Environmental Protection Agency.

^bAI, artificial intelligence and human-machine interaction; SWT, software technology; CPT, computer technology; BHW, basic hardware technology; CMT, communications technology; EDU, education and training; APP, applications and computational science; CAS, common applications support; CPI, computing infrastructure; CMI, communications infrastructure; ADM, National Coordination Office.

^cNC/GC, National Challenge/Grand Challenge; DCRT, Division of Computer Research and Technology (NIH); IAIMS, Integrated Academic Information Management System; NLM, National Library of Medicine; NCRR, National Center for Research Resources (NIH); NCI, National Cancer Institute (NIH); PDQ, Physician Data Query (NIH).

SOURCE: Data on agency budgets and program activities were extracted from the *FY 1995 Implementation Plan* prepared by the National Coordination Office (1994).

TABLE C.2 Actual FY 1993 and Requested FY 1995 HPCCI Budget (millions of dollars) Categorized by Discipline

Discipline	1993	1995	Percentage Change
Computer technology	63.66	107.00	68
Software technology	128.14	155.60	21
Communications technology	47.96	89.45	86
Computing infrastructure	165.60	204.72	23
Communications infrastructure	47.29	91.56	94
Applications and computational science	102.44	176.96	73
Common applications support	57.39	147.44	157
Artificial intelligence and human-machine interaction	36.85	57.65	56
Basic hardware technology	33.50	58.95	76
Education	44.66	64.54	45
Administration	0.40	0.70	75
TOTAL program	727.89	1154.56	59

Alternatively, the 11 discipline categories can be used to examine the balance between support for discipline-specific scientific research that uses high-performance computing and communications technologies and support for computer science research on new high-performance computing and communications technologies. Analysis of the FY 1995 HPCCI budget request shows that \$352 million (30 percent) would be invested in basic research in computer, software, and communications technologies; \$205 million (18 percent) in applied computer science research, artificial intelligence, and human-machine interaction; \$176 million (15 percent) in direct support of applications and computational science; and \$297 million (26 percent) in computing and communications infrastructure.

Commentary: Many Possibilities for Misinterpretation

The HPCCI has enjoyed a certain amount of political support and is growing even in a time of very tight federal budgets. The committee believes that this has created a "bandwagon" effect: the initiative has had its scope extended by the inclusion of some work not directly related to the HPCCI's goals, however valuable it may be, or work with broad relevance. The result has been a less than focused program.

For example, all high-performance computing and communications systems are built from electronics and depend directly on advancements in basic electronic technology. The ARPA Microsystems program activity, which constitutes the large majority of the basic hardware discipline, supports research in basic electronics technologies. This research will eventually benefit

the high-performance computing and communications technology base and help advance the nation's information infrastructure, but it could also be used in a wide variety of other contexts.

Another problem is the possible creation of false expectations about the extent to which the HPCCI could create the technology necessary for advancing the nation's information infrastructure. A large amount of work within the two applications disciplines is directed primarily toward the use of high-performance computing in solving certain scientific and agency mission problems. Only a part of this work, such as the creation of digital libraries, would apply directly to the goal of enhancing the nation's information infrastructure. Some of this work is directed at challenging computational science problems, which have excellent scientific impact but whose results are more easily justified as scientific results, rather than HPCCI results. Also, the HPCCI invests much more heavily in computing than in communications. Less than 16 percent of the FY 1995 request is for communications technology and infrastructure.

About one-third of the program is directed toward creating new technology directly applicable to advancing the information infrastructure. The growth in funding in these areas is offset by an unrelated decrease in research investment by industry, spurred in part by competitive changes in the computer and communications industries. As a result, the nation's total amount of research in high-performance computing and communications technologies is considerably less than, it appears, and in fact may be insufficient to maintain the strategic U.S. lead in these technologies or to support the rapid deployment of an enhanced information infrastructure.

NOTES

1. CPSMA (1994), p. 7; this report points out that labor-intensive, detailed disaggregation of published data may be the only way to understand how research program budgets are spent.

2. Amounts shown for FY 1995 are Executive Budget requests. At press time, agency appropriations had been made, but the involved agencies had not disaggregated the appropriations and reported the HPCCI portions to the National Coordination Office. A 2-year time period, FY 1993 to FY 1995, was used to help dampen any single-year jumps in level.

D**Current High Performance Computing and Communications Initiative Grand Challenge Activities**

Since its beginning, the High Performance Computing and Communications Initiative has included Grand Challenges, difficult scientific problems whose solution will yield new scientific understanding while simultaneously advancing high-performance computing and communications. The following list of current Grand Challenge activities is based on the FY 1994 and FY 1995 "Blue Books" and communications with National Coordination Office staff.

NATIONAL SCIENCE FOUNDATION*Aerospace*

- Coupled field problems

Computer Science

- Machine learning
- Parallel input/output (I/O) methods for I/O-intensive Grand Challenges

Environmental Modeling and Prediction

- Large-scale environmental modeling
- Adaptive coordination of results of predictive models with experimental observations
- Earthquake ground motion modeling in large basins
- High-performance computing for land cover dynamics
- Massively parallel simulation of large-scale, high-resolution ecosystem models

Molecular Biology and Biomedical Imaging

- Biomolecular design
- Imaging in biological research
- Advanced computational approaches to biomolecular modeling and structure determination
- Understanding human joint mechanics through advanced computational models

Product Design and Process Optimization

- High-capacity atomic-level simulations for the design of materials

Space Science

- Black hole binaries: coalescence and gravitational radiation
- Formation of galaxies and large-scale structure
- Radio synthesis imaging

DEPARTMENT OF ENERGY*Energy*

- Mathematical combustion modeling
- Quantum chromodynamics calculations
- Oil reservoir modeling
- Numerical Tokamak project

Environmental Monitoring and Prediction

- Computational chemistry (see Box D.1 for discussion)
- Global climate modeling
- Groundwater transport and remediation

Molecular Biology and Biomedical Imaging

- Computational structural biology

Product Design and Process Optimization

- First-principles simulation of materials properties

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

- Large-scale structure and galaxy formation
- Cosmology and accretion astrophysics
- Convective turbulence and mixing in astrophysics
- Solar activity and heliospheric dynamics

NATIONAL INSTITUTES OF HEALTH

- Molecular biology
- Biomedical imaging

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY

- Product design and process optimization

ENVIRONMENTAL PROTECTION ADMINISTRATION

- Linked air and water-quality modeling

NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION

- Climate change prediction and weather forecasting

BOX D.1 Computational Chemistry:
Applying High-Performance Computing to Scientific Problems

Chemists were among the earliest researchers to pursue the use of computers to extend understanding of their field. At the time of a 1956 Texas conference on quantum chemistry, electronic computers had developed to the stage that it was just feasible to program large scientific computations. However, these computers provided results too crude to be of interest to quantum chemists, even by the late 1960s. Nonetheless, based on this "progress," the conference passed a recommendation urging that more machines be at the disposal of university departments. Several groups at the University of Chicago, the Massachusetts Institute of Technology, and elsewhere pursued the goal of exploiting these new facilities, engaging in huge (for the time) research programs to compute the electronic wave functions for molecules constituted of two atoms from the first full row of the periodic chart.

By the early 1970s, significant progress had been made in the computation of molecular energies, wave functions, and dynamics of reacting systems and liquid structure by high-speed computers of the day, namely, the IBM 360/65, CDC 6600, and Univac 1108. Important work was also being done using semiempirical methods for systems with as many as 10 to 12 atoms. Reliable methods had been applied to the calculation of potential energy surfaces for $H + H_2$ and $F + H_2$ systems—methods that have been essential in the advancement of understanding molecular collisions. There had been semi-quantitative calculations of hydrogen bond strengths and protein conformation, but the facilities to carry out such work were available to only a small group of chemists, mostly on the staffs of national laboratories. The need to extend access to such facilities coupled with the new goal of bringing together people to work on software development and to attack important problems in chemistry led to the creation of the National Resource for Computation in Chemistry (NRCC) funded jointly by the National Science Foundation (NSF) and Department of Energy.

With the creation of the NSF supercomputer centers in the 1980s, chemists were able to pursue computational studies with requirements well beyond the capability of systems available otherwise even to leading research groups of the period. In addition to high-speed computation, the centers made accessible large amounts of disk storage and fostered large-scale use of high-speed data communication.

A major breakthrough of the early 1980s was the recognition by industry of the value of computational chemistry to the marketplace. Companies set up research groups that used computational chemistry software for electronic structure studies (e.g., Gaussian and CADPAC) and molecular mechanics simulations (e.g., AMBER and CHARMM), coupled with graphics platforms.

By the mid-1980s an industry had developed in modeling software focused on the chemical, pharmaceutical, and biotechnology industries. Large companies, such as Eli Lilly and Dupont, bought supercomputers to provide the capability to model complex molecules and processes important for their businesses.

One of the major directions for future work is the application of accurate quantum mechanical approaches to biological systems. This effort would complement the molecular mechanics method with selected calculations of higher accuracy to enable explication of important fine points. Areas where these efforts might be introduced are enzymatic reactions involving transition-metal centers and an array of catalytic processes.

The additional power provided by massively parallel computer systems is stimulating a push for higher accuracy and improved algorithms. Methods that have had impact for serial processors that were readily modified for vector systems often must undergo major modification or replacement for massively parallel processors. A major requirement for advancement is seamless scalability across systems of different size.

With the need for higher accuracy on massively parallel systems will likely come increased attention to Monte Carlo methods for quantum many-body systems. Quantum simulations are naturally parallel and are expected to be used increasingly on massively parallel computer systems.

E

Accomplishments of National Science Foundation Supercomputer Centers

INTRODUCTION

The National Science Foundation (NSF) Supercomputer Centers Program preceded the High Performance Computing and Communications Initiative but has become an integral and important part of it. The centers were established to provide access to high-performance computing—supercomputers and related resources—for the broad science and engineering research community. The program has evolved from one comprising independent, competitive, and duplicative computer centers to a cooperative activity, one that has been characterized as a MetaCenter.

In 1992 the four NSF supercomputer centers (Cornell Theory Center, National Center for Supercomputing Applications, Pittsburgh Supercomputer Center, and San Diego Supercomputer Center) formed a collaboration based on the concept of a national MetaCenter for computational science and engineering: a collection of intellectual and physical resources unlimited by geographical or institutional constraints. The centers' first mission was to provide a stable source of computer cycles for a large community of scientists and engineers. The primary objective was to help researchers throughout the country make effective use of the architecture or combination of architectures best suited to their work. Another objective was to educate and train students and researchers from academia and industry to use and test the limits of supercomputing in solving complex research problems. The best and most adventurous proposals for using an expensive and limited resource were sought.

In 1994, the scientific computing division of the National Center for Atmospheric Research joined the MetaCenter. In addition, the NSF established the MetaCenter Regional Affiliates program, under which other institutions could pursue projects of interest in collaboration with MetaCenter institutions. The MetaCenter thus became a unique resource and a laboratory for computer scientists and computational scientists working together on shared tasks.

IMPORTANT TECHNOLOGY ACCOMPLISHMENTS

Originally set up in 1985 to provide national access to traditional supercomputers, the NSF centers have evolved to a much larger mission. The centers now offer a wide variety of high-performance architectures from a large array of U.S. vendors. Today work at the centers is dominated by research efforts in software, in collaboration with computer scientists, focusing on operating systems, compilers, network control, mathematical libraries, and programming languages and environments.

Supercomputer Usage at NSF Centers

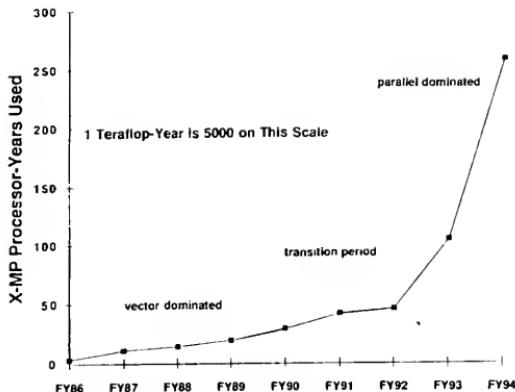


FIGURE E.1 Total historical usage of all high-performance computers in the NSF MetaCenter. This graph shows the total annual usage of all high-performance computers in MetaCenter facilities. Particularly striking is the growth since 1992, when microprocessors in various parallel configurations began to be employed. All usage has been converted to equivalent processor-years for a Cray Research Y-MP, the type of supercomputer that the NSF centers first installed in 1985-1986.

TABLE E.1 Supercomputer Usage at National Science Foundation Supercomputer Centers, 1986 to 1994

Fiscal Year	Active Users	Usage in Normalized CPU Hours ^a
1986	1,358	29,485
1987	3,326	95,752
1988	5,069	121,615
1989	5,975	165,950
1990	7,364	250,628
1991	7,887	361,037
1992	8,578	398,932
1993	7,730	910,088
1994	7,431	2,249,562

^aData prior to May 1990 include the John von Neumann Center.

Architectures and Vendors

The national research community has been offered access to a wide and continually updated set of high-performance architectures since the beginning of the NSF Supercomputer Centers Program in 1985. The types of architectures and number of vendors are now probably near an all-time high (Table E.2), allowing the science and engineering communities maximum choice in selecting a machine that matches their computational needs. A short list of architectures offered through the NSF centers program includes single and clustered high-performance workstations or workstation multiprocessors, minicomputers, graphics supercomputers, mainframes with or without attached processors or vector units, vector supercomputers, and single instruction multiple data (SIMD) and multiple instruction multiple data (MIMD) massively parallel processors.

TABLE E.2 Vector and Scalable Parallel High-Performance Computing Machines in the Four National Science Foundation Supercomputer Centers

FY 1991 7 Vector supercomputers
2 IBM 3090—6 processors each (new)
2 Cray YMP—8 processors each
Cray YMP—4 processors
Cray 2S—4 processors
Convex C240—4 processors
.....
FY 1991 4 Scalable parallel systems
2 TMC CM2—32,000 processors each
Intel iPSC/860—32 processors (new)
Alliant 2800—14 processors (new)
.....
FY 1992 6 Vector supercomputers
IBM ES9000/900—6 processors (upgrade)
2 Cray YMP—8 processors each
Cray YMP—4 processors
Cray 2S—4 processors
Convex C3880—8 processors (upgrade)
.....
FY 1992 11 Scalable parallel systems
2 TMC CM2—32,000 processors each
TMC CM5—512 processors
Intel iPSC/860—64 processors
nCUBE2—128 processors (new)
Kendall Square Research, KSR1—64 processors (new)
IBM PVS—32 processors (new)
Alliant 2800—14 processors
2 DEC Workstation Cluster (new)
IBM Workstation Cluster (new)
.....

continues

TABLE E.2—*continued*

FY 1993 Vector supercomputers
IBM ES9000/900—6 processors
Cray C90—16 processors (upgrade)
Cray YMP—8 processors
Cray YMP—4 processors
Cray 2S—4 processors
Convex C3880—8 processors
.....
FY 1993 14 Scalable parallel systems
2 TMC CM2—32,000 processors each
TMC CM5—512 processors
Intel Paragon—400 processors (upgrade)
nCUBE2—128 processors
KSR1—160 processors (upgrade)
IBM PVS—32 processors
IBM SP1—64 processors (new)
Cray T3D—32 processors (new)
2 DEC Workstation Cluster
IBM Workstation Cluster
HP Workstation Cluster (new)
MasPar 2—16,000 processors (new)
.....
FY 1994 5 Vector supercomputers
IBM ES9000/900—6 processors
Cray C90—16 processors
Cray C90—8 processors (upgrade)
Cray YMP—4 processors
Convex C3880—8 processors
.....
FY 1994 18 Scalable parallel systems
TMC CM2—32,000 processors
TMC CM2—8,000 processors
TMC CM5—512 processors
Intel Paragon—400 processors (upgrade)
Intel Paragon—28 processors (new)
nCUBE2—128 processors
KSR1—128 processors
IBM PVS—32 processors
IBM SP1—64 processors
IBM SP2—80 processors (new)
Cray T3D—512 processors
2 DEC Workstation Cluster
IBM Workstation Cluster
HP Workstation Cluster
MasPar 2—16,000 processors
Convex Exemplar—8 processors (new)
SGI Challenge—32 processors (new)

Current vendors whose top machines have been made available include IBM, DEC, Hewlett-Packard, Silicon Graphics Inc., Sun Microsystems, Cray Research, Convex Computer, Intel Supercomputer, Thinking Machines Corporation, and nCUBE, plus a number of companies no longer in existence, such as Alliant, Floating Point Systems, ETA, Kendall Square Research, Stellar, Ardent, and Stardent.

Access and New Architectures

In the 1960s, only a few universities had access to state-of-the-art supercomputers. By the early 1990s, some 15,000 researchers in over 200 universities had used one or more of the supercomputers in the NSF MetaCenter. This increased use led to new concepts and innovation:

- *Achieving production parallelism.* Cornell Theory Center became the first member center to achieve production parallelism on a vector supercomputer.
- *Migration to the UNIX operating system.* In 1987, the National Center for Supercomputing Applications (NCSA) became the first major supercomputer center to migrate its Cray supercomputer from CTSS to UNICOS, a UNIX-based operating system developed at Cray Research for its supercomputers.
- *Access to massively parallel computers.* NCSA introduced massively parallel computing (MPP) to the research community with the CM-2 in 1989, followed by the CM-5 in 1991. NCSA has worked closely with national users and the computer science community to create a wide range of 512-way parallel application codes that can in 1995 be moved to other large MPP architectures such as the T3D at PSC, the Intel Paragon at SDSC, or the IBM SP-2 at CTC.
- *Heterogeneous processors.* In 1991, the Pittsburgh Supercomputer Center was the first site to distribute code between a massively parallel machine (TMC-CM2) and a vector supercomputer (Cray Y-MP), linked by a high-speed channel (HiPPI).
- *Workstation clusters.* The NSF supercomputer centers were among the first to experiment with clusters of workstations as an alternative for scalable computing. The first work was done in the 1980s with loosely coupled clusters of Silicon Graphics Inc. workstations to create frames for scientific visualizations. With the introduction of the IBM RS6000, several centers moved to study tightly coupled networks and developed job control software. Clusters from DEC, Hewlett-Packard, and Sun Microsystems are now available as well.

Storage Technologies, File Format, and File Systems

With the vast increase in both simulation and observational data, the MetaCenter has worked a great deal on problems of storage technologies, with the greatest progress in software. The creation of a universal file format standard, a national file system with a single name space, and multivendor archiving software are some of the results of MetaCenter innovation and collaboration.

NSFNET and Networking

The 56-Kbps connection between the NSF supercomputer centers, established in 1986, was the beginning of the NSFNET. Based on the successful ARPANET and the TCP/IP protocol, the NSFNET rapidly grew to provide remote access to the NSF supercomputer centers by the creation of regional and campus connections to the backbone.

Although started by the pull from the high end, the NSFNET soon began to provide ubiquitous connectivity to the academic research community for electronic mail, file transport, and remote log-in, as well as supercomputer connectivity. As a result, the NSFNET backbone of 1995 has 3,000 times the bandwidth of the backbone of 1986. The centers have also developed prototypes for the high-performance local area networks that are needed to feed into the national backbone as well as the next generation of gigabit backbones.

Visualization and Virtual Reality

The NSF centers were instrumental in bringing the concepts and tools of scientific visualization to the research community in the 1980s. Center members developed new approaches to understanding large datasets, such as a three-dimensional grid of wind velocities and direction in a thunderstorm, by "visualizing" or creating an image from the data. This led scientists to consider visualization as an integral part of their computational tool kit. In addition, the centers worked closely with the preexisting computer graphics community, encouraging them to create new tools for scientists as well as for entertainment.

**Desktop Software, Connectivity,
and Collaboration Tools**

The history of the centers has overlapped greatly with the worldwide rise of the personal computer and workstation. It is, therefore, not surprising that software developers focused on creating easy-to-use software tools for desktop machines. These tools have had a major influence on the usefulness of supercomputer facilities to remotely located scientists and engineers, as have tools such as NCSA's telnet, which brought full TCP connectivity to researchers using IBM and Macintosh systems, significantly broadening the base of participation beyond UNIX users. Collaboration tools have provided the capability to carry on remote digital conferencing sessions between researchers. Both synchronous and asynchronous approaches have been explored.

Development of the nation's information infrastructure requires many software, computing, and communications resources that were not traditionally thought to be part of high-performance computing. In particular, tools need to be developed for organizing, locating, and navigating through information, a task that the NSF center staffs and their associated universities continue to address. Perhaps the most spectacular success has been the NCSA Mosaic, which in less than 18 months has become the Internet "browser of choice" for over a million users and has set off an exponential growth in the number of decentralized information providers. Monthly download rates from the NCSA site alone are consistently over 70,000.

ACCOMPLISHMENTS IN EDUCATION AND OUTREACH

Each of the supercomputer centers has developed educational and outreach programs targeted to a variety of constituencies: university researchers, graduate students, undergraduates,

educators at all levels, and K-12 students and teachers. Another aspect of outreach is the effort to identify and serve local and regional needs of government, schools, and communities. Activities range from the tours given at all MetaCenter installations through the hosting of visits by national, regional, and local officials and commissions, to full-scale partnerships. Table E.3 summarizes participation in these various activities.

TABLE E.3 Supercomputer Centers' Educational Activity Support Summary

Educational Activities	FY 1991	FY 1992	FY 1993
High school/K-12—Attendees	715	1,370	1,985
Research institutes—Attendees	262	377	390
Training courses and workshops—Attendees	1,700	2,400	2,100
Monthly newsletter circulation	234,986	247,692	165,176
Visitors	13,506	16,380	16,392

Researchers and Students

One- or two-day workshops offered by MetaCenter staff to researchers on-site and at associated institutions cover introductions to the computational environments, scientific visualization, and the optimization and parallelization of scientific code. In addition, special workshops have been offered throughout the MetaCenter on the use and extension of computational and visualization techniques specific to various disciplines.

MetaCenter institutions have contributed to the research projects of hundreds of graduate students through the provision of fellowships or similar appointments, stipends, access to resources, and relationships with MetaCenter researchers. Programs providing research experiences for undergraduates bring in students to work for a summer or a school semester or quarter on specific projects devised by MetaCenter researchers and/or faculty advisors. In many instances such projects have resulted in presentations at meetings and publications.

K-12 Educators and Students

Training of high school teachers and curriculum development are among the many MetaCenter educational efforts. Several programs have been initiated, such as ChemViz to help students understand abstract chemistry concepts; a visualization workshop at *Supercomputing '93*; and SuperQuest, a program involving MetaCenter sites that brings teams of teachers and students from selected high schools to summer institutes to develop computational and visualization projects that they then work on throughout the following year.

Broad Outreach

Outreach is also accomplished by the publications programs of the MetaCenter, the production of scientific videos and/or multimedia CD-ROMs, and a collaborative program for

maintaining a lively and informative presence on World Wide Web servers, which make information on the MetaCenter's programs easily accessible over the nation's information infrastructure.

A number of interactive simulation programs are now being tested in classrooms across the country and around the world. Students can change initial conditions and watch a simulation evolve as the parameter space is explored. The educational programs of the MetaCenter made available to high schools around the country demonstrate the power of the nation's information infrastructure to provide new educational resources.

SCIENTIFIC COMPUTATION AND INDUSTRIAL DEVELOPMENT

Partnerships between the MetaCenter and industry are collaborations with major and large industrial firms, as well as small companies and venture start-ups. Most of these partnerships exist because MetaCenter expertise has been essential to the introduction of new ways of using the resources of supercomputing: the algorithms, visualization routines, and engineering codes are being combined in ways that result in such advances as high-end rapid prototyping of new products.

Commercialization of the software developed at the MetaCenter is being undertaken by a number of companies. For example, NCSA telnet has been commercialized by Intercon, and Spyglass has commercialized NCSA desktop imaging tools, as well as its Mosaic program. CERFnet, a California wide area network for Internet access, has pioneered in supplying access to library holdings and other large databases, and DISCOS/UniTree, a mass storage system, is in use at more than 20 major computer sites. A new molecular modeling system, called Sculpt, developed at the San Diego Supercomputing Center, is being commercialized by a new company, Interactive Simulations. Sculpt enables "drag-and-drop" molecular modeling in real time while preserving minimum-energy constraints: its output was featured on a May 1994 cover of *Science*.

IMPORTANT SCIENCE AND ENGINEERING ACCOMPLISHMENTS

Selected areas and problems, summarized below, indicate the range of projects currently being undertaken by nearly 8,000 researchers at over 200 universities and dozens of corporations and the span of disciplines now using this new tool.

Quantum Physics and Materials Science

The great disparity between nuclear, atomic, or molecular scales and macroscopic material scales implies that vast computing resources are needed to attempt to predict the characteristics of bulk matter from fundamental laws of physics. Since the beginning of the NSF centers program, researchers in this area have been among the most frequent users of supercomputers. Materials scientists have often been among the first to try out new architectures that promise higher computational speeds.

Listed below are some examples of research areas important to the study of properties of bulk matter in extreme conditions, such as occur in nuclear collisions, the early universe, or the core of Jupiter; new materials such as nanotubes and high-temperature superconductors; and more practical materials used today such as magnetic material and glass.

- Phase transitions in quantum chromodynamics
- Phase transitions of solid hydrogen
- New nanomaterials predictions

- Theory of high-temperature superconductors
- Magnetic materials

Biology and Medicine

Living creatures exhibit some of the greatest complexity found in nature. Therefore, supercomputers have made possible unprecedented opportunities to explore these complexities based on the fundamental advances made in biological research of the last 50 years. These activities include using the data from x-ray crystallography to study the molecular structure of macromolecules; learning how to use artificial intelligence to fold polypeptide chains, determined from genetic sequencing, into three-dimensional proteins; and determining the function of proteins by studying their dynamic properties.

New fields of computational science, such as molecular neuroscience, are being enabled by academic access to MetaCenter computing and visualization resources and staff. Corporations are using supercomputers and advanced visualization techniques in collaboration with the NSF MetaCenter to create new drugs to fight human diseases such as asthma. New insights into economically valuable bioproducts are being gained, for instance, by combining molecular and medical imaging techniques to create "virtual spiders" that can be dissected digitally to understand the production of silk. Finally, high-performance computers are becoming powerful enough to enable researchers to program mathematical models of realistic organ dynamics, such as the human heart. Examples of projects include the following:

- Crystallography
- Artificial intelligence and protein folding
- Protein kinase solution
- Molecular neuroscience—serotonin
- Molecular neuroscience—acetylcholinesterase
- Kinking DNA
- Antibody-antigen docking
- Tuning biomolecules to fight asthma
- Virtual spiders and artificial silk

Engineering

Man-made devices have become so complex that researchers in both academia and industry have turned to supercomputers in order to be able to analyze and modify accurate models in ways that complement traditional experimental methods. High-performance computers enable academic engineers to study the brittleness of new types of steel, improve bone transplants, or reduce the drag of flows over surfaces using riblets. Industrial partners of the individual supercomputer centers within the MetaCenter are using advanced computational facilities to improve industrial processes such as in metal forming. Better consumer products, such as leakproof diapers or more efficient airplanes, are being designed. Even state agencies are able to use the MetaCenter facilities to improve traffic safety or find better ways to use recycled materials. Some 70 corporations have taken advantage of the MetaCenter industrial programs to improve their competitiveness.

Examples of engineering-related problems include the following:

- Heart modeling
- Ultrahigh-strength steels

- Continuous casting of steel
- Beverage-can design
- Designing a leakproof diaper
- Bone transplant bioengineering
- Improving performance with riblets
- Designing better aircraft
- Crash-testing street signs

Earth Sciences and the Environment

The resources of the NSF MetaCenter are being used to compute and visualize the complexity of the natural world around us, from the motions of Earth's convective mantle to air pollution levels in southern California. The U.S. Army is working with academics to determine how to practice tank maneuvers without endangering the breeding habits of the sage grouse. Pollution is a difficult coupling of chemical reactions and flow dynamics that must be understood in detail if corrective measures are to be efficacious. High-performance computers also act as time machines, allowing for faster-than-real-time computation of severe storms. Finally, to improve global weather or climate forecasts, supercomputers allow researchers to study the physics of such critical processes as mixing at the air-ocean interface. Among the related problems being addressed are the following:

- Detoxification of ground water
- Storm modeling and forecasting
- Los Angeles smog
- Upper-ocean mixing
- Simulating climate using distributed supercomputers

Planetary Sciences, Astronomy, and Cosmology

As was evident in the recent impact of Comet Shoemaker-Levy 9 with Jupiter, observatories on Earth and in space have become intimately linked. Supercomputers are being integrated into observational facilities, like the Grand Challenge Berkeley Illinois Maryland Association millimeter observatory, and into observational programs such as the ones that have led to the discovery of new millisecond pulsars or the first extrasolar-system planet.

The ability of numerical methods to solve even the most complex of fundamental physical laws, such as Einstein's equations of general relativity, is increasing understanding of the dynamics of strong-field events, such as the collision of black holes. In perhaps the grandest-scale challenge possible, the universe itself is a subject of investigation by several Grand Challenge teams using resources of the MetaCenter to discover how the large-scale structures in the universe evolved from nearly perfect homogeneity at the time of the formation of the microwave background.

- Comet collision with Jupiter
- Discovery of first extrasolar system planet
- Pulsar searching and discovery
- Black hole collision dynamics
- Cosmological simulations

F
**Individuals Providing Briefings
 to the Committee**

Duane Adams
 Advanced Research Projects Agency

John deFerrari
 General Accounting Office

William Andahazy
 House Armed Services Committee

Howard Frank
 Advanced Research Projects Agency

Dan Anderson
 Ford Motor Company

Norm Gjostein
 Ford Motor Company

Forest Baskett
 Silicon Graphics Inc.

Wayne Hamann
 Ford Motor Company

Gordon Bell
 Computer science consultant

Peter Highnam
 Schlumberger Austin Research Inc.

James Burrows
 National Institute of Standards and
 Technology

Lee Holcomb
 National Aeronautics and Space
 Administration

John Cavallini
 Department of Energy

Kenneth Huebner
 Ford Motor Company

Melvyn Ciment
 National Science Foundation

Elizabeth Johnston
 General Accounting Office

George Cotter
 National Security Agency

Anita Jones
 Department of Defense

Craig Davis
 Ford Motor Company

Thomas Kalil
 National Economic Council

Hassan A. Dayem
 Los Alamos National Laboratory

Donald Lindberg
 National Library of Medicine and
 National Coordination Office

Ed McGaughan
Office of U.S. Senator Jeff Bingaman

Michael Nelson
Executive Office of the President

Steve Nelson
Cray Research Inc.

Merrill Patrick
National Science Foundation

Richard Radtke
Ford Motor Company

Justin Rattner
Intel Supercomputer Systems Division

Victor H. Reis
U.S. Department of Energy

Guy L. Steele, Jr.
Thinking Machines Corporation

Paul Strassmann
Independent consultant/writer

Sing Tang
Ford Motor Company

John Toole
Advanced Research Projects Agency

Steven J. Wallach
CONVEX Computer Corporation

Philip Webre
Congressional Budget Office

James Wilson
House Science, Space, and Technology
Committee

Charles Wu
Ford Motor Company

Mr. SCHIFF. I am going to ask all members to try to stay roughly within 5 minutes for questioning. If at the conclusion of every member's opportunity to question, if someone has a burning issue that they want to get into, I would of course recognize them, and that will apply to me as the chair also.

I have two matters I would like to bring up. One is I want to say at the outset that my view of the—first of all, I guess I should say case in point about the advancement of computers, Dr. Sutherland, your point. I bought a computer system for my congressional office in 1991. I am told today that it is so out-of-date that they cannot find parts for it. They have to jury-rig it when something goes wrong, and I have to go get another one.

Also, one quick war story. There was a group of us in Chicago where I grew up—this is all part of my 5 minutes, colleagues. I want to stress that.

[Laughter.]

Mr. SCHIFF. But there was a group that where we grew up, we played chess together in high school and college. One fellow upon graduating college got a job doing computer work at Northwestern University. He wanted to see if a computer could be taught to play chess. So, on his own time he began to secretly teach his computer how to play chess, and when they found out about it at Northwestern, he thought for sure he was going to be fired. They gave him a full-time job of teaching his computer to play chess. We have seen some of the results of that both in chess and in other aspects of computers.

I have two matters I would like to bring up. One is I want to say that I have a very strongly positive view toward the idea of Government-private industry cooperation. In fact, I believe that in that regard I am probably at some point beyond the view of the majority of the Members of the House who I think would cut that short of where I would promote cooperation between Government and business because I see that happening around the world, and I think it is important to the United States in this new era of global competition.

Dr. Jones, you made a statement, though, that even caught my attention even with where I am coming from and that is industry is not in the business of developing new technologies. Industry is in the business of making profit. Well, can one not argue that developing new technology is making profit, that if one rings the bell on an invention and patents it and legally controls the market for a period of time, they make a very substantial profit on that? So, is there really not a question here of would business not do more if they were really motivated to do so, if they really saw the value in it? I would yield to you.

Dr. JONES. Certainly industry does make exciting technology inventions. The point I was trying to make is that when you see these very rapid product turn cycles of 18 months, those are very incremental steps that are being taken. They are effective but they are incremental steps being taken.

It is the case in this country today that the vast majority of research and development done by U.S. industry in this arena and others has a short-term view, short-term being 18 months, 3 years, sometimes 5 years, but it is quite a small minority portion. I think

Dr. Sutherland gave you some of the reasons why it is an abiding Government interest investing in the long term, and it is hard for corporations to do. Again, I think it is a partnership.

Mr. SCHIFF. Mr. Toole, do you want to further respond?

Mr. TOOLE. I might just add to that and I think you will hear some subsequent testimony in the second panel that some of the definitions we use from research, development, long-term, short-term certainly vary in some of this. I think I strongly support what Dr. Jones said because we have had several interactions with industry leaders.

For example, we sponsored 2 weeks ago, invited in basically the chief technical officer level to come in and speak to a Government group of people of where their futures were, and I think it becomes very enlightening in terms of where their, quote, research and development monies are actually being applied for.

I think when we are speaking from the Government side of long term, we want to do fundamental things down the road 5 to 15 years that can have an impact and be able to capitalize on the research and development that is being done by both industry and academia and the Government.

Mr. SCHIFF. Mr. Sutherland, any further response? I know you testified on this subject.

Dr. SUTHERLAND. I would make two comments here. One is if you look at the turnover in the shares of American companies, I asked, and it turns out the New York Stock Exchange 25 years ago took about 7 years before all of the shares had traded hands. Today it takes less than 2 years for all of the shares on the New York Stock Exchange to change hands. So, the owners of American business are in fact not owners very long, and so their interest in the long term is less than you might expect.

What industry has in fact proven itself very good at doing is improving the product base that it now makes, and it appears to me, based on 15 years experience in the venture capital business, that the way U.S. industry faces new tasks is in fact to start new companies. I work for one that is only 13 years old. It filled a need in the computer business that the existing companies perhaps knew about and for whatever reasons chose not to enter this market, leaving a market opportunity for a new company to in fact come in and offer products which are very valuable.

Mr. SCHIFF. I am going to conclude my time here just by making this observation. Again, I am a strong believer in partnership between industry and Government, and so I think we are all on the same side. But the testimony I am hearing indicates to me that there is a niche for free enterprise, that the company that might boldly invest in long-range research might ultimately reap the greatest rewards.

With that, I would like to recognize Mr. Geren for 5 minutes.

Mr. GEREN. Thank you, Mr. Chairman. I want to thank our panel for their very interesting and thought provoking testimony.

Dr. Sutherland, your point about the turnover in personnel in companies. I thought that was a very interesting insight. I had never considered that as one of the impediments to investment in R&D in companies, but that makes sense to me.

The National Academy of Sciences report suggests that the HPCC Program has subsumed too much of the Federal computing research budget, and the report also recommends that the program focus be tightened so that it does not include all or even most of the computing and communications research activities supported by such agencies as NSF and ARPA.

Could you all comment on the Academy's recommendations? I think, Dr. Sutherland, you talk about or you alluded to the need to avoid too tight a focus I guess because you do not know really where basic research is going to take you. You do not know where you are going to end up, and that is part of the value of basic research. At one point, I guess I would like you all to address is there a danger in trying to focus it too much, that you would lose the genius of the undertaking?

Dr. SUTHERLAND. I think there is a good thing about HPCCI which is it has fostered this cooperation and collaboration between the various independent agencies. I think another good thing about it is it has not usurped, it has avoided usurping the agencies' mission responsibilities.

I think the weakness that we have seen has been a tendency to use HPCCI as an umbrella to say, oh, this is marginally HPCCI related. Let us put it in the HPCCI program because that is a convenient umbrella to provide a funding resource. What we recommend is that the basic programs in the agencies are of vital importance to the Nation and must be continued.

The HPCCI spirit of collaboration which has happened is a wonderful and good thing, and John Toole's office I think provides an important forum for the folk who are responsible for those programs to meet and collaborate.

What I think needs to be avoided is the use of the HPCCI or any initiative of that sort as an umbrella for only marginally related activities that are involved in the area.

Dr. JONES. As I had said earlier, I believe that the activities in the High Performance Computing and Communications Program really are core activities for the agencies involved, certainly the lead agencies, ARPA, NSF, DOE, and NASA. These basic fundamental investments in information technology, progress in information technology are crucial to each of the agencies, and by its lights, by its agency mission decides how to make its investment.

Many of the things that are cited as being important high performance computing results are, in fact, not included in what today is called the program. One of the books that you have in front of you is mission successes based on high performance computing. It is in the Department of Defense. These are all applications of high performance computing. It does not talk about the research there but, in fact, an advance that was enabled by having high performance computing, the hardware, parallel software, and tools, and in fact is not included in the funding of the program. It really is production kind of work.

I think all of us at this table have said in one way or another that there is a core long-term research program in fundamental computation, the hardware, the architecture, the software, the networks, and that that is a crucial national investment. So, we are

in fact trying to follow the advice of the National Research Council committee along those lines.

Mr. GEREN. Mr. Toole, before you answer, let me just ask you to be more specific. Do you plan to carry out the recommendation 12 in the report which suggests the need for review of the projects to identify long-term research areas that should be independent of the program initiatives?

Mr. TOOLE. Specifically to your question, we have an interagency review process where we review each of the individual programs together and collaboratively and critique it. We are in the process of actually writing up those recommendations as an internal document that we are going through that.

So, the direct answer to your question is, yes, we are looking very carefully at what aspect of it, and I would like to pose the HPCC initiative and its subsequent investments to be all long-term investments coupled to some of these mission needs, so when we look at programs like ASCI that we have made reference to, which are very important users, a small piece of that may be R&D fitting into the activities of what they need to meet their mission, but certainly not the bulk of that particular kind of a program. Likewise, we have many, many downstream users that are involved with that.

I would like to go back to your original question, if I may, sir.

I fully support what Dr. Sutherland said. I think in this interagency process—and that is why I need the priority and support of the agencies. They are my customers as well as things that make it happen. It is their management expertise internally and together that really makes the difference. We continually, virtually every year, have to cut certain kinds of activities.

So, I would like to be sure that we do not have any kind of marginal efforts whatsoever in the criteria based on the Sutherland report, based on our own internal organizations, as well as many of the very serious budget issues that are arising in today's budget process across the different agencies from reorganizations to downsizings to different ways. It comes up over and over again that the HPCC technologies become a priority within those agencies to be able to effect the kinds of changes and cost effectiveness that they need.

So, I wanted to assure you that I certainly take the challenge to be able to look at the program and across the program to determine and help determine, coupled with the agencies themselves who are senior managers and people in our own organizational process, to ensure the highest quality activity for the long term.

Mr. GEREN. Do you agree with Dr. Sutherland's advice about what your role should be, and that is primarily as a coordinator?

Mr. TOOLE. I absolutely do. We have found, for example, things that have evolved through the Internet, things that have evolved successfully with our technology in academia. The very best way that can be achieved is empowerment through innovation and not to have a centrally managed, overall arching kind of a program. So, part of my role and responsibility here is to facilitate and encourage in such strategic areas coupled with the agencies that they can execute and manage and encourage the innovation to come forward

from the bottom as well as through the processes that we have already got in place.

Mr. GEREN. Thank you. Thank you, Mr. Chairman.

Mr. EHLERS. [presiding] Thank you, Mr. Geren.

I have several questions.

Incidentally, I apologize on behalf of Chairman Schiff. He had to leave for approximately 30 to 45 minutes for another committee meeting, and he will return later. In the meantime, I will fill in as Vice Chair of the Full Committee.

I, first of all, want to thank you for your presentations. I have very few questions about the presentations themselves. I found them very clear and enlightening and appreciate being brought up to date on some of these issues.

I do have an overall question. Dr. Jones and Mr. Toole, you both indicated you did not think reauthorization was required, and I am a bit puzzled by that. Why would you not want the Congress to reauthorize this particular project?

Dr. JONES. I think that information technology is so important and so fundamental to where this Nation is going and needs to go in science and in industry that it is and should be just the core program that this Government funds.

I know for our participation in this program, we do not think of investment in the long-term research in information technology as an initiative but as a decadal investment. Most of the defense investment has been through the Advanced Projects Research Agency, and I look to the initiative as a direction from Congress that said these agencies are going their own way funding information technology research. The Nation would be better served if they came together and coordinated, and I think the initiative has accomplished that, but this is such a fundamental, broad-based investment that I believe it ought to be considered a core investment by the agencies that depend on the progress in this area.

Mr. EHLERS. Mr. Toole?

Mr. TOOLE. Mr. Ehlers, I am not sure I can add much more to that. We believe we are trying to institutionalize this high priority investment in information communications. In terms of how we want to go forward with reauthorization, we think strong support at the individual agency levels is the way, coupled with strong support of the interagency process that we have put before you. We will gladly work with you in any way, shape, or form that we think is the best strategy to make this happen, as we discussed.

Mr. EHLERS. I am a little puzzled perhaps because I am a newcomer to the Congress, but it seems to me the agencies were going their own way and Congress came along and said, no, we want you to work together. So, the initiative was developed, and now you are saying it is too important to be left to the Congress. The agencies will make their own decisions because this is so important.

Dr. JONES. Oh, excuse me, sir.

Mr. EHLERS. I thought you might want to respond to that.

[Laughter.]

Dr. JONES. Let us see. I think Congress was exactly right, but the message is different. The coordination was an exceedingly good idea. It is a cost-effective idea and I think we had better, more highly leveraged, serendipitous even results. But it is now a way

of life for the agencies, and we have not only taken it on board and followed it but found out that it worked and it is now an integral part of the process. I think the challenge is to do this as well in some other arenas.

Mr. EHLERS. It seems to me what you are really saying is you are not opposed to reauthorization, but you are opposed to continuing this as an initiative, that if reauthorization takes place, it should recognize it as an ongoing entity rather than something that is new and has to be experimented with.

Mr. TOOLE. Precisely.

Mr. EHLERS. All right.

Dr. Sutherland, I was fascinated with some of your examples and also your analysis of why industry does not do research. I think all three points are accurate.

I do, however, have to reinforce my pet peeve on your third point and that is the shareholder actions. Frankly, I feel rather strongly because I have been in the research field for a number of years. I really think that many corporations are shirking their responsibility and I also think that many shareholders are irresponsible in taking such a short view of the value of their investment. I agree with you that new companies form to fill the void, but there is something tragic when the larger corporations become such a bureaucracy that they cannot entertain new ideas and make investments that will yield a profit. And I know from friends in the field that this has happened a number of times.

You do not necessarily have to respond to that, but I certainly agree with the point. I just have to get in some condemnation of American industry. It seems to me a century ago the name of the game was entrepreneurship, taking risks, and getting a return on investment. We seemed to have lost a lot of that.

If you wish to comment, you may.

Dr. SUTHERLAND. I have served on the boards of public companies, as well as a venture capitalist on the boards of private companies, and what I observe is that in privately held companies, the owners of the company behave in a way responsible for the long-term stewardship of the company. In publicly held companies, many of the shareholders are forbidden to participate in the management by law. Consequently, although they hold the shares, they do not behave like owners, and they have little interest in the long-term health of the company because they can leave it at any moment.

I think the fault here rests not with the management of companies but with the ownership of them. The fact that the New York Stock Exchange turns over in less than 2 years—I work at Sun Microsystems. We have about 100 million shares outstanding and we often trade a million shares a day, which means that the entire stock base of the company turns over in a very short period. Now, you can argue that underneath that trading there is a long-term holding. Nevertheless, it is the rapid traders that set the stock price, and the board, believe me, is responsive to that.

Mr. EHLERS. I agree with that, and as I said, I think the problem is shareholder irresponsibility. You may not wish to use that word but I use it without hesitation.

One specific question, Dr. Sutherland, just a matter of curiosity to me. You list in your report, which you referred to in your testimony and I took a quick look at it, the results of Government-sponsored research and you listed the mouse and windows as coming from Government-sponsored research. That is a new one to me.

Dr. SUTHERLAND. Well, the mouse was invented by Douglas Englebart at the Stanford Research Institute. I happen to be working at ARPA at the time and I was the responsible officer for the contract. So, I can assure you that as Government-supported.

Mr. EHLERS. I was not aware of that.

[Laughter.]

Mr. EHLERS. So, Apple Computer did not do it after all or Xerox, either one.

Thank you very much.

Next we go to Mr. Doyle.

Mr. DOYLE. Thank you very much. I want to start by thanking our Chairman for holding today's hearing and welcome our distinguished witnesses. I think this is a very important program and I am glad we are having these hearings today because I think it is a program that absolutely should be reauthorized.

I am one of the new kids on the block, being a freshman Member of Congress, and it is my understanding that when this program started, it had strong bipartisan support. But there are many of us in the 104th Congress that are concerned about the future of this program. Already we have seen a \$35 million reduction in NASA, and this Committee acted to cut NASA's HPCC budget during consideration of H.R. 2043 without the benefit of hearings.

I think it is important we have these hearings today so that members understand not only how important the program is, but the positive impact the program is having on the American economy and on academia. I wonder if you could maybe elaborate a little bit on some examples, also of how this program has helped Government operations to be more cost-effective.

A couple of examples in the Department of Defense. It is my understanding we have been able to develop imaging systems that operate at a much greater cost effectiveness, and that the National Security Agency also makes great use of supercomputing in its code-breaking operations.

So, as we take a look at this program and some of us are worried that this may attacked as an example of corporate welfare—we are hearing that word bandied about a lot in the 104th Congress—I think it would be important for members and for the record to reflect how this program helps Government be more cost effective, and I wonder if the panel could expand on some of the examples I have given or some others.

Dr. JONES. Mr. Doyle, you have given some excellent examples. A couple more come to mind.

In the Department of Defense, each member of the armed forces receives many weeks of training every year, and we are turning to simulated-based, computation-based training for initial teaching of procedures and skills so that people are already up at a threshold before you have to spend a fairly large number of dollars to move them to an actual live training range. So that is an area that we

see some production of costs coming from high performance computing and indeed it does take high performance computing.

I cited weather prediction where, when you can give exact and reliable predictions, people believe them and they leave and lives are saved. They do board up businesses and protect the property better, and so we actually see reduction in losses.

This goes I think on and on. There are a number of examples I think you can translate to cost effectiveness and cost reduction some of the examples that you see in what is called the Blue Book that you have before you from the high performance computing.

Mr. TOOLE. Mr. Doyle, I might add a couple and certainly some of our other people that we have invited to come could certainly add to it.

But, for example, the virtual laboratory, being able to tie laboratories and remote facilities together is becoming a way in which many of the technologies from communications, distributed systems, the ability to put software environments together is being really strongly considered and perhaps the only way in which some of these downsizing activities within the Federal Government can actually be achieved. The education and training that Dr. Jones hinted at was another one.

Shared resources in general I think are becoming a goal between the different agencies where you have got very large scale machines, medium scale, and other activities that can actually be used.

So, the fundamental things that we are really working with here become the implementation means by which a lot of this can happen. Obviously the world wide web itself, making Government information more readily available across the Internet, and putting information on line I think is having just a phenomenally fundamental impact on our society, our Government, and the way we are exposing more and more information. It is somewhat experimental in some cases but we are providing valuable information to the citizen, the people, and directing more directly which I think translates in real terms into dollars.

Mr. DOYLE. Would it be possible to submit some of these examples, to submit detailed examples for the record of some of these programs, and also maybe to submit for the record the amount of dollars being saved as a result of these programs so that we can have that reflected in the record?

Mr. TOOLE. I would be happy to do so although I might point out that dollars saved could be somewhat controversial because we are a long-term research program. So, therefore, I think we can cite longer examples that have evolved over time and similar kinds of investments we are making now, but I just caution the fact in terms of the dollars saved. We would be honored to do this for the record.

Mr. DOYLE. Thank you very much.

Is my time up, Mr. Chairman?

Mr. EHLERS. The gentleman's time has expired.

Mr. DOYLE. Thank you, Mr. Chairman.

Mr. EHLERS. Next the gentleman from Minnesota, Mr. Gutknecht.

Mr. GUTKNECHT. Thank you, Mr. Chairman. This is a very interesting hearing. When I looked at the group of people who were going to be testifying today, I almost wished that Dr. Seymour Cray from Minnesota was going to be here.

In fact, a couple of years ago I had a chance to tour the University of Minnesota Supercomputer Center, and I found it fascinating. But one of the things that sort of bothers me is that there is a lot that they are doing there that they cannot tell you about. It is either national defense or it is work that is being done by private corporations that they want to keep under very tight wraps.

I was wondering, Mr. Toole, could you talk a little bit about—and you have talked a little bit about successes. When we try and sell this to some of our colleagues in the Congress. They want to know what are we getting for this. We heard about the mouse. Are there any other things we can talk about?

Mr. TOOLE. Well, I think there is a long list that we would be more than happy to provide you. I think Dr. Sutherland hit the majority of them, really on some of the key areas over the long term that have made a major, major difference. I think some of the investments in terms of the national HPCC Program has been completely open. In other words, we have downstream users that perhaps may be classified or cannot talk about what that research is.

But the program—in fact, this was a very important part of actually doing things at a national level—allows the technology and the technology base to be completely open and available. I think based on your constituency or the Congress, that is exactly what we hope we are providing in some of our testimony and some of these reports because I think they give you a very long opportunity and list of key elements that are really making a difference in today's world and where we think we need to be in the future.

The web network technology I think is becoming central to everyone's environment or will become central to everyone's environment, as well as the computing resources in very distributed ways. So, if you look at where you want to go next in mobile and distributed worlds, applying that to organizations that can actually save and make money for themselves, whether it be internally or externally, that is some of the fundamental things that I think we are trying to look at.

Mr. GUTKNECHT. I want to get back to another point, and I think we have talked about this a little bit already in that our society, whether we are talking about shareholders or taxpayers or just people in our society today, is becoming incredibly risk averse. America was founded by people who took risks. The people who got on boats and came over here basically said you got a king, you got your taxes, I am sorry, I am out of here. I am going to America, the land of opportunity. Yet, in the last number of years, we have sort of moved away from that.

I want to come back to the example that Dr. Sutherland talked about with the mouse where, in effect, with the ARPA grant I assume someone—and I forgot the name—developed the mouse. We also believe in the United States about risk and reward. So, the United States Federal Government—and I want to hear more about this story—took some risks and helped finance this deal. Did we get anything back for it? Did we get any reward?

Dr. SUTHERLAND. Well, that is an interesting case in point. The mouse was invented sometime in the 1960's, so it was amazingly early. At the same time what the group at SRI was working on was what is now called hypertext. The idea was that you could, in fact, have a piece of text and it would have links within it, so if there was a part of a paragraph that you were interested in more information for, you could hit it with the mouse and basically go to the other place.

Now, in those days the computing equipment was not really good enough to do that. It was only marginally possible to make demonstrations of capability. Nobody in his right mind would think of that as a commercial activity. Well, today it is a commonplace kind of thing.

The remarkable thing about that story is how long it takes from when it is a gleam in somebody's eye making a demonstration of capability and the steps that it takes to go from there to commercial practice.

I will tell you the history. I believe this is more or less correct. The mouse was picked up at Xerox Park with Xerox corporate dollars, and the system of on-line computers called the Alto was made which explored both networking and workstation kinds of activities and had the mouse as a part of the routine affair. Some folks from Apple came and had a tour of the Xerox Palo Alto Research Center and saw this thing and said, gee, that would make a wonderful commercial product. They in fact were able to commercialize it when Xerox, for reasons that are not clear to me, failed to do so.

I think the remarkable story about American industry is that all too often existing industries have a success story going. They have a product base which is the right product base, and they are very good at improving that product base. The existing companies seem not to be so good at recognizing brand new opportunities and capitalizing on them. So, you see the MacIntosh appearing from a brand new organization that was hardly heard of and sort of taking the world by storm, and then other organizations gradually pick up those ideas, seeing that they are successful, and the technology begins to pervade the industry.

So, the question of what did we get back out of the mouse, well, I do not know how many mice are in use in the Federal Government, but I expect there is a fair number. The efficiency that is providing in the operation of Government is quite remarkable, to say nothing of the taxes and the payroll taxes that are being paid by the companies that were formed to do the work, and to say nothing of the exports that they generate, and to say nothing of the efficiency of the whole society that comes about because personal computers are now a part of the fabric of our Nation. It is hard to say what would have happened if that invention had not happened, and it is hard to say, hey, if Englebart had not done that, somebody else would have. It is hard to make those cases.

But the history is fairly clear, that it is often the Government-funded academic institutions—Stanford Research Institute is a not-for-profit institution near Stanford University—that think far enough in the future to say one day hypertext will be important and so we need a device like this to use.

What is going on today is we are looking forward into the first part of the 21st century when we are going to have 1,000 times more computing than we have now, and who knows what that is going to do for us? Who knows what the bright young people are going to come up with as the right idea that makes that thousand-fold improvement in computing be the strength of our Nation?

Now, I wanted to respond a little to an earlier point about where do the HPCCI dollars go. More than half of it, I understand from John Toole, goes into academic institutions, and of the remaining half, about a half of the remaining half, or a quarter of the total, goes into not-for-profit and Government laboratories. It is only a quarter of the HPCCI budget that is going to industrial activities. I think that is quite proper, that the long-term innovation in the United States comes mainly from people who take a long-term view and those people are mainly in our academic institutions. We have a marvelous situation, a very strong academic research program in this country, and I think it is one of the major assets that we have for the future.

Dr. JONES. May I follow up?

I think one of the direct products of the high performance computing program has been the students that are educated. They populate the new companies of this day and the next day. They are key individuals in making the technical decisions of how information technology is applied. So, they are more than just the inventors of the new ideas. The program gives you more than the ideas. It gives you the highest expertise talent in the country, and it is very important to us to sustain that small pool but very highly leveraged pool of people.

Mr. EHLERS. The gentleman's time is expired.

I just want to mention to Dr. Sutherland I have been a confirmed MacIntosh user for almost a decade, and I thank you for making that original grant to SRI.

[Laughter.]

Mr. EHLERS. We would like to move on to the next panel as soon as possible. Mr. Baker or Mr. Weldon, do you have questions? Mr. Baker?

Mr. BAKER. With that admonition, Mr. Chairman, let me just turn it around. You have been saying that Government has to continue supporting research in the computing area, and I agree. What would Government do to get out of your way? What can we do to the tax structure? What can we do to the prohibitions against you working together so that we can become stronger as an exporter? The FDA prohibits pharmaceutical companies from coming to market in their own lifetime. What is Government doing in your area that we could improve?

With that, I would like to pass the time on to my colleague here from Florida and let them think about it because obviously we are doing everything perfectly and we do not need to improve.

Mr. EHLERS. The gentleman from Florida, Mr. Weldon.

Mr. WELDON. Sorry for my late arrival. I think this is a very important hearing nonetheless.

The only question I would like to address to the panel—and I do not know if anybody has posed this already. I am under the impression there are 12 Federal agencies that are currently involved

with the program and it is up for reauthorization. Is that a problem having 12 different Federal agencies involved with this act just from a management perspective?

Dr. JONES. We think it is a strength. One of the strengths of this program is that the agencies followed congressional direction which said go work together and leverage what each other are doing, and today we are fielding some programs that I think no one agency could field, but because it is coordination through the National Co-ordination Office, we have some stronger activity than would otherwise exist. So, to the contrary.

Mr. TOOLE. I think just to follow up on that, yes, it is indeed a challenge, but I think it is a challenge worth taking and going forward with. The results have been quite evident as a result to avoid duplication, to open people's eyes in terms of what other agencies are doing, to get people in the trenches to understand what digital library technologies could be together, for example, within the agencies. It is through this dialogue and discussion and debate that really things start to come to the surface to have an impact. So, I fully support that, although I am sort of in a position where that is my job, but I really strongly believe in that.

Mr. WELDON. Thank you. I yield back the balance of my time.

Mr. EHLERS. Thank you.

We are graced by the presence of Mr. Walker from Pennsylvania, Chairman of the full Committee. Do you have any comments or questions?

The CHAIRMAN. Yes, thank you, Mr. Chairman. I do appreciate that and I am just being briefed here. I realize this may be covering old ground but I understand that the testimony earlier was that you do not feel as though there is a need to go to a new authorizing act and that instead the authorizations can be taken as a part of each of the individual agency authorizations. Is that correct?

Dr. JONES. Yes.

The CHAIRMAN. Tell me about the impact that that is going to have on some of the designated high performance computing centers. Does that mean that in order to do their budgets, they are going to have to seek money from each of the agencies and therefore be in a position of kind of getting their money in a rather divided stream, or are they still going to be able to seek one kind of appropriation through NSF and then have NSF find the money from the various agencies? How do you envision that working?

Dr. JONES. Let us see. Today we coordinate the different agency programs. The High Performance Computing and Communications Program is not run with a single, centralized budget. So, I do not think this will change materially the issues and the decision process for funding centers or research programs.

Mr. TOOLE. I would like to confirm that, Mr. Chairman. In addition, I think our plan is to continue to produce the kind of documents that are cross-cutting and have the visibility of the coherence across the different agencies that participate, and we seek your guidance and the Committee's guidance on how best we can achieve that and improve on that in the future.

The CHAIRMAN. And it will not affect the ability at all of universities and others to access the network to lack an authorization for the specific program in your view?

Mr. TOOLE. In my view we expect strong support from the individual agencies and commitments, and we believe we have gotten that in their respective agencies. So, we want to go forward with the kind of commitment that we currently have in the program. I think you may have missed some of our testimony, but basically 52 percent of the overall program today goes to academia, and I think we want that to be as strong if not stronger in the future.

The CHAIRMAN. This may enhance the ability of universities to participate in your view rather than simply having the Federal Government make the determinations and therefore the allocations?

Mr. TOOLE. I think in this particular program, which is really a program for the long term, we really have opened and have a lot of participation with academia. I think in the next panel you will hear some testimony by some people from academia, and I think that question probably is more appropriately posed to them in terms of their participation. But I am a very strong supporter of academic research and their ability to leverage that to beat the agency corporately with the missions that we have.

The CHAIRMAN. Thank you, Mr. Chairman.

Mr. EHLERS. Thank you for your attendance.

The Vice Chair of the Committee, the gentleman from Texas, had one more quick question.

Mr. GEREN. If we were not to reauthorize it and each agency were to just operate independently in furthering what is already in place, I would just like you all to comment on the National Academy's report where they think there should be more restrictive criteria. It should be applied to selecting Grand Challenge research projects. In other words, only projects should be selected that contribute substantially to advancing high performance computing.

How would these decisions be made on what are the Grand Challenges that these efforts should be coordinated towards? How would you provide that kind of guidance in an even less coordinated operation than you have currently?

Mr. TOOLE. I really would like to hope—both elicit your support as well as the agencies' support—and I think this is kind of where it comes down to—for a very strong interagency process. So, I do not believe that we should see decreased interagency activity. If anything we are seeing that this is now becoming part of the institutionalization within the agencies.

So, to answer your specific question, I think that is exactly what we would get involved with in the yearly and bi-yearly strategic planning processes and budgeting processes across the agencies and in conjunction with the Congress to determine precisely what areas are needed, as well as what areas are going to be the most fruitful for the long-term research needs of the Nation.

Dr. JONES. Two comments. One is I do not think that this will lead to less coordination among the agencies because they found it so fruitful in the past.

Second, there is a tension between investing the next dollar in the particular mission agency application and against investing

that dollar in long-term research in, say, software for scalable parallel applications. The agencies today make both choices with some balance. That tension will remain, and I think the Committee was just cautioning that the agencies not let the single mission application drive out the enabling fundamental research. I think one can always argue that the balance is not quite right, but I think the agencies are taking a responsible approach to that today.

Mr. GEREN. And you think that even without reauthorization, as you move forward, it would be possible to get the agencies to come together and agree on what the Grand Challenges are and coordinate among all the agencies that those are what we pursue?

Dr. JONES. Coordination will continue because it has been so successful, yes.

Mr. GEREN. I guess it is not the coordination I am getting to, but how you decide what the Grand Challenges are. The Academy is saying that they do not think that you have done a good job of selecting Grand Challenges in the past, that perhaps it has been a little bit too unfocused. They are saying we need to pick some Grand Challenges, and that is their term not mine.

Dr. SUTHERLAND. May I offer a little clarification on that since I was part of that Academy panel?

Mr. GEREN. Yes.

Dr. SUTHERLAND. I think what we looked at was that the Grand Challenge program had been chosen initially as in part an exercise for very high performance parallel computing, and that the emphasis in choosing the projects to be done, the challenges to be done, assessed their exercising of the computing part quite heavily and said if they did not do that a lot, they probably were not worth doing.

What we are saying is in some sense there has been success. We have in fact exercised the computing and demonstrated enough of the parallel computing that in fact you can see that it is valuable, and that now in selecting such programs, they should be selected more for their agency value. The DOE should be selecting Grand Challenges that are energy related. DOD should be selecting challenges that are defense related, and NSF should be selecting challenges that are science related. The science, defense, and energy portions of those programs should be emphasized over and above the computing portion. There should be some computing ingredient, but I do not think we said that they were badly chosen. We just said that now the emphasis on how to select them can be changed a little bit.

And I think it fits well with what is being proposed here in terms of saying that this information processing stuff is so valuable to the Nation, it is not just an initiative. It is a continuing base of activity that we have to have.

The CHAIRMAN. Will the gentleman yield?

Mr. GEREN. I would be glad to.

The CHAIRMAN. I just want to follow up on that. What I am hearing you say then is that what you want to do is switch all the program toward the applications, toward the mission oriented concerns. My concern is that we maintain the fundamental underlying research. That explanation leads me to think that the lack of reauthorization is aimed at assuring that these agencies, as they spend

their money, are going to spend the money directed toward only those things that relate to their mission. What I want is an assurance that we do not back away from doing the fundamental underlying research that presses the envelope on high performance computing.

Dr. JONES. It is crucial to continue to do that, and I think the slide, the billboard that is over against the wall there that Dr. Sutherland used, highlights the productivity of the fundamental Government investment. It is the basis for the discovery or the invention of time sharing and many other fundamental things that you are alluding to.

The CHAIRMAN. But is there a danger as you shift the emphasis more toward the agencies that they fall into exactly the pattern of deciding that the mission is more important than the fundamental research? That is really my question.

Mr. TOOLE. Quite frankly, Mr. Chairman, there is always a danger of that. Let me explain a little bit how we intend to deal with that.

We want this HPCCI initiative continuation in the long term for information technology to be the long-term component piece of it. When we are speaking of applications here, I do not believe we are only speaking of applications just inside this particular initiative. In fact, some of the recommendations in the Sutherland report point to some of those things that may be more in the downstream user base. For example, this DOE ASCI Program that we have used as an example here is not part of the HPCC Program today. Yet, it is absolutely clear that it is dependent upon the results of investments we have made in our particular program.

I hope this will continue, and I think this is exactly what we want to work for, that we have many, many programs across the Federal Government, not just the R&D programs like this one which really needs a focus on that front end, but have enough pilot initiatives associated with mission drivers to make it relevant and important for the actual application areas. I think we are seeing that type of activity, and that is where we would like to be able to assure in the next phase of this program.

The CHAIRMAN. I thank the gentleman for yielding.

Mr. EHLERS. The gentleman's time has expired.

I want to thank the panel again for their excellent testimony. We certainly appreciate your coming here and not only testifying but answering the questions. Thank you very much.

We will now ask the second panel to step forward. Thank you very much for your participation today. We will immediately begin with Dr. Ingram.

STATEMENT OF DR. JOHN INGRAM, RESEARCH FELLOW, SCHLUMBERGER, AUSTIN, TEXAS

Dr. INGRAM. Thank you, Mr. Chairman. I appreciate being invited here particularly to discuss a subject which is very dear to my heart and to the industry I work for.

My comments are from a different standpoint of industry; that is, we are users of high performance computing. For me high performance computing is at the beginning. It is just coming of age. The physical capabilities of computers, memory, and communica-

tions are approaching the levels necessary to solve important real problems.

Were the problems of the past not important or real? Certainly they were important. The successes have been dramatic. New aircraft designs have profited greatly from the simulation process. Seismic processing has opened new fields for oil exploration. Dramatic results have come from biochemical applications. They were as real as the scientists and engineers who designed them could make them.

These applications and others like them, however, are in transition. They are moving from a one and two-dimensional phenomenological world to three or more dimensions with quantitative results. This is creating a paradigm shift that can move supercomputing into the mainstream of industrial life. There are, however, some significant barriers to overcome.

I would like to emphasize that industry is concerned with reality. Basically its products must be real.

Computing is a tool for industry, a tool for discovery, for decision making, for efficiency, and competitive advantage, and above all a tool for change. In the discovery stage, high performance computers have been a great asset. Research has yielded new understanding and ultimately new concepts for products and services. These have then had to go through the painful process of prototyping, industrialization, manufacturing, and implementation. It is in this downstream part of the process that the move to supercomputing has been disappointingly slow. The reasons for this give indications for our future actions.

Industry is by nature opportunistic. We heard that a bit this morning. Its primary goal, after all, is to make money and few indulge in research blindly in the hope of finding an interesting commercial application.

I would like to add an off comment there. I work for a company that actually has engaged in long-term research for a very long time. It has made a great deal of profit from its long-term research. I will give you two examples which are not in here.

One is that there has been for many years the feeling that a measure of permeability in oil reservoirs could be an extremely profitable product. We still do not have a direct measure of permeability. That research project has been going on for about 30 years now.

The development of a nuclear accelerator to produce neutron bursts and non-chemical sources has been going on now for 20 years, and we still do not have it. I hope some day we will have it. If we do, by the way, we will make a lot of money from it.

While the need to see a product at the end of the road necessarily limits our horizons, it also provides a connection to reality. Research and development money is carefully analyzed for risk versus return. The current state of supercomputing makes its inclusion on the critical path for product or service introduction doubtful. First of all, the instability of the manufacturers themselves raises questions concerning long-term product evolution. Perhaps more important, the lack of standards and software for these machines means that massive development efforts may be lost when the next generation of machines arrives.

I do not want to imply by that that I do not feel like that it is worth doing. It is very much worth doing. It just requires caution.

This is even truer for supercomputers than for the rest of the rather chaotic computer business. Over the years, these machines have been developed by a process not unlike an automobile manufacturer making bigger and bigger engines and neglecting the steering, transmission, and suspension. In the end, the car goes very fast in a straight line but would be dangerous drive on mountain roads. Industrialization of a product is rarely without bumps and curves.

In order to understand where we stand in the application of high performance computing to commercial problems, it is useful to have a look at a real example that has already benefitted greatly from them. oil exploration.

Much has been said about supercomputers as time machines to get scientific results quickly rather than waiting for the evolution of simpler machines. Supercomputers are also time machines for industry but in a different sense. Time lines such as time to market and turnaround time are of primary importance. Often how long it takes to get a result to a client will determine whether it is useful to him or not. In the early stages of exploration, an oil company may well commission a three-dimensional seismic survey. This may involve weeks of data acquisition and yield up to a thousand billion bytes of data.

An aside there again to give perspective how much data that is, even with current recording techniques, the best way to transmit that data in the highest band width possible is to use a 747.

Extracting the structure of the subsurface can then require months of computation. In order to shorten the time to make a decision on whether or not to develop the oil field, a time machine is invaluable. Shortening this process has been one of the dramatic successes of parallel computing. The certainty of finding oil has also greatly increased with the change from two-dimensional to three-dimensional surveys and has repaid the cost of supercomputing many times over. The future, however, lies in going well beyond.

If I am saying nothing else, I am saying we are really at the beginning. We are three orders of magnitude from being able to solve the problems that we would like to be able to solve.

When the oil man decides to develop the field, he is again faced with considerable uncertainty. Wishing to minimize his return, he would like to instrument the reservoir and to plan the recovery process.

A comment there. The oil reservoir is probably the largest, non-instrumented resource in the world.

He would like that at each stage the choice of well locations, flow rates, injection rates, and other parameters are optimized for total recovery and speedy return on investment. His decisions will depend on getting quantitative responses to key questions. The word "quantitative" is very important. The characterization of the reservoir must be statistical and each well he drills gives him new information. These are then combined into a massive computer simulation that is constantly updating, including new 3D seismic surveys to track fluid interfaces and multiple computer runs to define

most likely scenarios. Now the time scale he is dealing with is years.

If he succeeds, it may increase the yield of hydrocarbons by as much as 10 to 20 percent. For a major reservoir, this can be billions of barrels of oil equivalents. Is this an adequate incentive to undertake the simulation? Only if the computer speed, memory capacity, and communications capability permit valid results in time to control the evolution of the reservoir.

Of course, there are other pieces missing from this picture besides supercomputing resources. One such is the quantitative characterization of statistically inhomogeneous reservoirs, and another is the accurate prediction of multi-phase fluid migration. While it is reasonable to expect the petroleum industry to provide the answers to these fundamental questions through its own research, it must in turn be able to depend on commercially available systems for processing and communications.

This case illustrates a fundamental shift in the use of high performance computing. With the availability of three and four-dimensional results that truly represent nature and business conditions, the motivation to use the most powerful computing systems available will be strong. The investment in understanding and deploying the measurement systems in creating the simulation environment and putting in place the communications facilities adequate for monitoring is enormous.

In addition, a great deal more information enters into the operation of the asset than just knowledge of the subsurface. Economic conditions may lead to the construction of electricity generating facilities in the vicinity of the reservoir, and pipeline availability may influence timing. Management of such assets represents the confluence of information technology of ever-increasing complexity. There is little hope of accomplishing such an ambitious undertaking without truly professional computing and communications environments.

The paradigm shift to three dimensions and quantitative results is happening throughout science and industry. We have finally admitted that the world is three-dimensional and that realism, albeit expensive, is paramount. At the same time, the economic climate for supercomputers has changed, and one question is. Will there be an evolving, rational family of such machines in the future and will there be an adequate software environment to support it?

There is every reason to believe that when the dust settles and the relations between Government, academia, and industry are clear, a healthy high performance computing sector will emerge. The free market system can provide all of the financial incentive needed for this. There is little doubt that industry will step up to the challenge if the risk is of reasonable proportion. For example, the shift from two to three dimensional seismic surveys increase the cost by factors of 20 to 50 times. Yet, there is more seismic activity today than ever before.

What will not happen without simulation is the research in high performance computing environments. We must make up for the sins of the past and build the software and communications structures to support the hardware. This should be applications-driven research. It should not be done in vacuum but in close contact with

the scientists and engineers who will implement and use it. How can we bring this about?

One cause of the high performance computing software problem is the vastly different cultures and operating modes of the United States Government, academia, and industry. Bringing these together is a difficult task but it must be done. It is fundamental that the research institutions understand the uses of their results.

The simplest and surest way to promote industry involvement in the research programs of Government and academia is to move people. This is considerably less difficult if the research is clearly of interest to the industry in question. Consortia are another means of attacking the problem. In any case, the cooperative programs are the surest way to achieve practical implementations of research.

The research funded by HPCC is of prime importance to American industry. Each industry has its own agenda and only through the efforts of HPCC will the research in key areas of information technology and high performance computing take place. A result of the information technology explosion is that the world is shrinking. In order for American industry to succeed, it must compete in the world marketplace. There is a growing awareness that America has excellent opportunities outside our boundaries but only as a member of the world community. We have limited funds for research and there is good work going on elsewhere. We need first to decide our priorities in the areas in which we must be the best. We must learn to leverage these funds through joint efforts with our worldwide partners. Perhaps it is time to consider global out-sourcing.

[The prepared statement of Dr. Ingram follows:]

High Performance Computing in Industry

by John D. Ingram

High Performance Computing is coming of age. The physical capabilities of computers, memory, and communications are approaching the levels necessary to solve important real problems. Were the problems of the past not important or real? Certainly they were important. The successes have been dramatic. New aircraft designs have profited greatly from the simulation process. Seismic processing has opened new fields for oil exploration. Dramatic results have come from Biochemical applications. They were as real as the scientists and engineers who designed them could make them. These applications and others like them are, however, in transition. They are moving from a one or two dimensional phenomenological world to three or more dimensions with quantitative results. This is creating a paradigm shift that can move supercomputing into the mainstream of industrial life. There are, however, some significant barriers to overcome.

Computing is a tool for industry; a tool for discovery, for decision making, for efficiency and competitive advantage, and above all a tool for change. In the discovery stage, high performance computers have been a great asset. Research has yielded new understanding and ultimately new concepts for products and services. These have then had to go through the painful process of prototyping, industrialization, manufacturing, and implementation. It is in this downstream part of the process that the move to supercomputing has been disappointingly slow. The reasons for this give indications for our future action.

Industry is by nature opportunistic. Its primary goal, after all, is to make money and few indulge in research blindly in the hope of finding an interesting commercial application. While the need to see a product at the end of the road necessarily limits our horizons it also provides a connection to reality. Research and Development money is carefully analyzed for risk vs return. The current state of supercomputing makes its inclusion on the critical path for product or service introduction doubtful. First of all the instability of the manufacturers themselves raises questions concerning long term product evolution. Perhaps more important, the lack of standards

and software for these machines means that massive development efforts may be lost when the next generation of machines arrives. This is even truer for supercomputers than for the rest of the rather chaotic computer business. Over the years these machines have been developed by a process not unlike an automobile manufacturer making bigger and bigger engines and neglecting the steering, transmission, and suspension. In the end the car goes very fast in a straight line but would be dangerous to drive on mountain roads. Industrialization of a product is rarely without bumps and curves. In order to understand where we stand in the application of High Performance Computing to commercial problems it is useful to have a look at a real example that has already benefited greatly from them: Oil Exploration.

Much has been said about supercomputers as time machines to get scientific results quickly rather than waiting for the evolution of simpler machines. Supercomputers are also time machines for industry but in a different sense. Time lines such as "time to market" and "turn around time" are of primary importance. Often how long it takes to get a result to a client will determine whether it is useful to him or not. In the early stages of exploration an Oil Company may well commission a Three Dimensional Seismic Survey. This may involve weeks of data acquisition and yield up to a thousand billion bytes of data. Extracting the structure of the subsurface can then require months of computation. In order to shorten the time to make a decision on whether or not to develop the oil field, a "time machine" is invaluable. Shortening this process has been one of the dramatic successes of Parallel Computing. The certainty of finding oil has also greatly increased with the change from two dimensional to three dimensional surveys and has repaid the cost of supercomputing many times over. The future, however, lies in going well beyond.

When the Oilman decides to develop the field, he is again faced with considerable uncertainty. Wishing to maximize his return, he would like to instrument the reservoir and plan the recovery process so that at each stage the choice of well locations, flow rates, injection rates, and other parameters are optimized both for total recovery and speedy return on investment. His decisions will depend on getting quantitative responses to key questions. The characterization of the reservoir must be statistical and each well he drills gives him new information. These are then combined into a massive computer simulation that is constantly updated including new 3D Seismic

Surveys to track fluid interfaces and multiple computer runs to define most likely scenarios. Now the time scale he is dealing with is years. If he succeeds it may increase the yield of hydrocarbons by as much as ten to twenty percent. For a major reservoir this can be billions of barrels of oil equivalents. Is this an adequate incentive to undertake the simulation? Only if the computer speed, memory capacity, and communications capability permit valid results in time to control the evolution of the reservoir. Of course there are other pieces missing from this picture besides supercomputing resources. One such is the quantitative characterization of statistically inhomogeneous reservoirs and another is the accurate prediction of multi-phase fluid migration. While it is reasonable to expect the Petroleum Industry to provide the answers to these fundamental questions through its own research, it must in turn be able to depend on commercially available systems for processing and communications capabilities.

This case illustrates a fundamental shift in the use of High Performance Computing. With the availability of three and four dimensional results that truly represent nature and business conditions, the motivation to use the most powerful computing systems available will be strong. The investment in understanding and deploying the measurement systems, in creating the simulation environment, and in putting in place the communications facilities adequate for monitoring is enormous. In addition, a great deal more information enters into the operation of the asset than just knowledge of the subsurface. Economic considerations may lead to the construction of electricity generating facilities in the vicinity of the reservoir and pipeline availability may influence timing. Management of such assets represents the confluence of Information Technology of ever increasing complexity. There is little hope of accomplishing such an ambitious undertaking without truly professional computing and communications environments.

The paradigm shift to three dimensions and quantitative results is happening throughout Science and Industry. We have finally admitted that the world is three dimensional and that realism, albeit expensive, is paramount. At the same time the economic climate for supercomputers has changed and one question is: Will there be an evolving, rational family of such machines in the future and will there be an adequate software environment to support it? There is every reason to believe that when the dust settles and the relations between Government, Academia, and Industry are clear, a

healthy High Performance Computing Sector will emerge. The Free Market System can provide all of the financial incentive needed for this. There is little doubt that industry will step up to the challenge if the risk is of reasonable proportions. For example, the shift from two to three dimensional Seismic Surveys increased the cost by factors of twenty to fifty times and yet there is more seismic activity today than ever before. What will not happen without stimulation is the research in High Performance Computing environments. We must make up for the sins of the past and build the software and communications structures to support the hardware. This should be Applications Driven research. It should not be done in vacuum but in close contact with the Scientists and Engineers who will implement and use it. How can we bring this about?

One cause of the High Performance Computing Software problem is the vastly different cultures and operating modes of the United States Government, Academia, and Industry. Bringing these together is a difficult task but it must be done. It is fundamental that the research institutions understand the uses of their results. The simplest and surest way to promote industry involvement in the research programs of Government and Academia is to move people. This is considerably less difficult if the research is clearly of interest to the industry in question. Consortia are another means of attacking the problem. In any case cooperative programs are the surest way to achieve practical implementations of research.

The research funded by HPCC is of prime importance to American Industry. Each industry has its own agenda and only through the efforts of HPCC will research in key areas of information technology and high performance computing take place. A result of the Information Technology explosion is that the world is shrinking. In order for American Industry to succeed it must compete in the World Market Place. There is a growing awareness that America has excellent opportunities outside our boundaries but only as a member of the World community. We have limited funds for research and there is good work going on elsewhere. We need first to decide our priorities and the areas in which we must be the best. We must learn to leverage these funds through joint efforts with our world wide partners. Perhaps it is time to consider Global Out Sourcing.

MR. EHLERS. Thank you very much.

We have been notified of a vote, in fact a series of votes. We have at most 5 minutes before we have to leave to go vote. Dr. Lazowska, if you can complete your testimony in 5 minutes, we will proceed with you, and then we will have to recess for approximately 15 minutes.

STATEMENT OF DR. EDWARD D. LAZOWSKA, PROFESSOR OF COMPUTER SCIENCE AND ENGINEERING, UNIVERSITY OF WASHINGTON, AND BOARD MEMBER, COMPUTING RESEARCH ASSOCIATION

Dr. LAZOWSKA. I will do my best.

Mr. Chairman and members of the Subcommittee, thanks for the opportunity to testify. My name is Ed Lazowska. I head the Department of Computer Science and Engineering at the University of Washington. I was a member of the NRC panel on which Ivan Sutherland just testified. I also serve on the Board of Directors of the Computing Research Association, and it is the 200 members of the CRA who I represent today. These are academic departments of computer science and computer engineering and industrial research labs in computer science and engineering where the Nation's cutting edge research and education in these fields takes place.

The Computing Research Association strongly endorses the findings of the NRC study which were discussed by Ivan Sutherland. Let me just recap five of these briefly.

First, there has been really mind-boggling progress in information technology which pervades most aspects of our lives.

Secondly, the nation that leads in information technology has enormous competitive advantages.

Third, America owns that leadership today and it is thanks to a really complex interaction and interplay between Government, academia, and industry.

Fourth, the Government's role is crystal clear. Industry can afford to look ahead only a few years, but as a nation we can and must invest for the long-term future. By and large, this work takes place in academic institutions with Government support.

Fifth, the HPCC Program has been a major success and that is in a variety of ways. The emphasis on high performance is appropriate. Cutting edge information technology is indeed a time machine.

Let me be concrete in the vein of an example that Dr. Sutherland gave.

I was an undergraduate in the late 1960's. The sine qua non of computer systems then was the IBM system 360. It had just been introduced in the middle 1960's and in fact architected in large measure by Fred Brooks, who was the co-chair of the NRC panel. A high end 360, which my undergraduate institution had, would occupy a room probably half again as large as this one. There would be a whole floor underneath it of offices for the people who kept it going by working on it or praying around it, depending on how things were going.

[Laughter.]

Dr. LAZOWSKA. This MacIntosh, which Mr. Ehlers will recognize—it is now 3 years old, so hopelessly out-of-date—dominates a high end 360 in every respect. It has more memory. It executes more instructions per second. It has larger disk storage than a typical university computing center of the late 1960's. So, that is a measure of the hardware progress, a concrete measure.

Now, as Ivan testified, one reason that we know how to use this computer—how to do hypertext, what you see is what you get editing—is that people in the 1960's used large building size mainframe computers as time machines. I was an undergraduate at Brown University and a gentleman there, Professor Andy van Dam, along with Englebart, was one of the developers of this hypertext technology. Year after year from midnight till 8:00 in the morning, I and my colleagues were stashed away in the computer center using this mainframe as a single-user text editor. That is what this time machine means. That is how we learned how to build the software systems and the applications that make it possible for you to get something out of this computer that you can carry around in your briefcase.

There are other aspects HPCC that have been a huge success as well. The emphasis on parallel computing. There is much more to be done but its viability is clear, and many important science and engineering problems have been solved. And the interagency co-ordination and cooperation is working very well.

So, we believe that continued authorization of the HPCC program elements is essential to the Nation. That means continued funding in these critical research areas and continued strengthening of the interagency process. And here are five reasons why.

First, the HPCC Program is the Nation's research and education program in information technology. HPCC is a coordinated multi-agency initiative that supports nearly all of the Nation's fundamental research and graduate education in information technology. It is important to recognize the breadth of that program. HPCC is much more than supercomputer centers. It is much more than the highest performance systems, although these systems are indeed time machines, and the path from cutting edge to desk top is really direct. HPCC is systems, software, networking, human resources, technology and applications for the Nation's information infrastructure.

I firmly believe that information technology is the Nation's future, and the HPCC Program at present is our Nation's research and education program in information technology.

Now, there were 12 of us on the NRC panel and this probably means that there are 12 different subtle degradations of interpretation of the report, but I think Ivan and I would agree that in response to Mr. Geren's question and Mr. Ehlers' questions, the committee was attempting to make clear that information technology should be supported independent of any special initiative. That is at the crux of it, and one of you gentleman said that. I think it is important for Congress to figure out what the best way to accomplish that is.

HPCC does have a very broad interpretation. This is appropriate because of the insertion of the information infrastructure technology and applications component. What is critical is that the Na-

tion's future be assured through continued funding of work in information technology.

The second point I want to make is that the interagency coordination really has been a model success. To maximize the likelihood of success in risky endeavors requires multiple agencies and multiple approaches. There has to be coordination. There cannot be tight management. There are lots of specific examples of terrific coordination in HPCC. The supercomputer centers, the gigabit testbeds and networking, the digital library initiative, and indeed the CIC strategic plan.

Mr. EHLERS. If you will pardon me.

Dr. LAZOWSKA. Yes.

Mr. EHLERS. We are at the last minute to leave to go vote. So, I hate to interrupt you, but I will have to. We will have to recess for approximately 15 minutes or as soon as we can get back from the second vote, and we will pick up right where you are now.

Dr. LAZOWSKA. Fine.

Mr. EHLERS. My apologies but our schedule is dictated by the bells and the clocks. Thank you.

[Recess.]

Mr. EHLERS. I would like to ask everyone to take their seats.

Mr. SCHIFF. I believe, Dr. Lazowska—I may be mispronouncing that. Am I?

Dr. LAZOWSKA. You are doing as well as my family does.

[Laughter.]

Mr. SCHIFF. [resumes chair] My family originally comes from Lodz which is spelled L-o-d-z, so I am accustomed to those errors.

Anyway, I believe we were in the middle of your testimony and please continue.

Dr. LAZOWSKA. Thanks.

I was in the process of enumerating a set of five reasons why I felt that continued authorization of the HPCC Program elements was critically important, that is, continued funding of this research, continued strengthening of the interagency process.

The first thing I pointed out was that the HPCC Program really is the Nation's research and education program in information technology.

The second is that interagency coordination has really been a model success, and I had just finished enumerating a set of examples of this.

The third point I wanted to make is that the HPCC Program has proven to be appropriately flexible and adaptable, and it is important to recognize that in research this is essential. Fundamental research is by its nature unpredictable.

An analogy is that when Lewis and Clark were exploring the country's geographical frontiers, they had a strategic objective, but things did not always turn out the way they planned them. There were false starts. There were changes in directions and emphasis.

This is also the case with HPCC as we explore the country's technological frontiers. The name has stayed the same, the program has evolved and adapted. The focus on software has increased dramatically. The focus on communications has increased dramatically. The entire information infrastructure technology and applications component was added to the program. It has a dramatically

increased emphasis on research issues related to horizontal scale, that is, ubiquity, in addition to research aspects related to vertical scale. So, HPCC has proven its ability to adapt.

Fourth, a strategic plan for the future exists. Much is done, much remains to do. You heard from Anita Jones and others about the CIC strategic plan. It included six strategic focus areas which I think are a good road map for the future: global scale information infrastructure technologies, high performance/scalable systems, high confidence systems, virtual environments, user centered interfaces and tools, and human resources and education. The strategic planning effort of the CIC was a major effort that I want to endorse strongly, as well as endorsing the directions that I identified.

The fifth point I would like to make—and again, you heard much about this before from Dr. Sutherland and others—is that the role of universities and the Federal Government is critical. The historical track record is clear. Over the course of many decades, federally supported university research has played a critical role in every aspect of information technology. You saw Dr. Sutherland's chart from the NRC study.

I serve on a six-person technical advisory board for Microsoft. Over the past 5 years, Microsoft has discovered that in order to create new markets, it needed new technologies in areas such as data compression, encryption, networking, 3D computer graphics, operating systems, statistical decision theory, and so forth. As demonstrated by the Brooks-Sutherland report, without America's research universities, these and other technologies would not be available to spur our world leadership.

Universities look to the future. The HPCC Program has been a huge success in enabling this through the time machine phenomenon. It is important to emphasize that university research carried out under a program avoids picking winners and losers. The purpose of publicly funded research in high technology fields is to advance knowledge and create new opportunities that industry can exploit in the medium and long term. The purpose is not to determine how the markets should develop. That is not what we do.

Universities transfer technology in two ways. They transfer ideas by granting patent licenses (by placing concepts in the public domain) and they transfer people (students and faculty leave to join or form companies).

Close industry-university interactions facilitate this tech transfer as well as the exchange of insights about long-term strategic directions. This is a pattern of innovation in technology transfer, fluid interaction between academia and industry, that has made America the world leader in information technology, and it will help us maintain this critical lead.

To summarize, we believe continued authorization of the HPCC Program elements, that is, continued funding of these research areas and continued strengthening of the interagency process, is essential to the Nation. My personal view is that reauthorization of the HPCC initiative would be helpful but it is perhaps not necessary. What is necessary to the Nation is that this work and this coordination go forward because it is our future. The name is HPCC. The program is America's program in advanced information technology.

Any authorization should be flexible in its approach, focus on fundamental research in a broad range of areas, and allowing adaptation as new targets of opportunity appear. The research areas described in the strategic plan are an excellent framework.

As you have heard from many witnesses before, we need to increase the focus on software both for computation and communications. We need to keep in mind that applications are important research drivers and paradigm shifters. Applications such as those covered in the user-centered interfaces and tools component of the strategic plan.

All of us here today understand the constraints under which the Subcommittee and the Congress is working. I think all of us also believe that the Federal investment in information technology research has been incredibly small compared to the payoffs to the Nation.

Thanks for the opportunity to testify.

[The prepared statement and attachments of Dr. Lazowska follow:]

Statement of
Edward D. Lazowska
Chair, Department of Computer Science & Engineering
University of Washington
and
Chair, Government Affairs Committee
Computing Research Association
U.S. House of Representatives Committee on Science
Subcommittee on Basic Research
Hearing on the High Performance Computing and Communications Act
October 31, 1995

Mr. Chairman and members of the Subcommittee, thank you for the opportunity to testify on the subject of the reauthorization of the High Performance Computing and Communications Act. My name is Ed Lazowska. I head the Department of Computer Science & Engineering at the University of Washington. In addition, I serve on the Board of Directors of the Computing Research Association (CRA) and head its Government Affairs Committee. I am here in that capacity today, representing the nearly 200 industrial research laboratories and academic departments in computer science and computer engineering that are members of CRA and that perform most of the nation's cutting-edge research and graduate education in the critical fields of computer science and computer engineering.

I have two other affiliations that provide perspective that I will bring to bear on this testimony (although I am not here representing these organizations). First, I am a member of the six-person Technical Advisory Board for Microsoft Research. This position affords me a firsthand view of the interplay between this leading information technology company and the national research enterprise. Second, I am a co-author of the recent National Research Council report *Evolving the*

High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure, which you have heard discussed already today by Dr. Ivan Sutherland.

I am pleased that the Subcommittee has decided to hold these hearings at this time. The Act is now about four years old and will expire in one more year. It is an appropriate time to ask what the program has achieved and where we should go from here. As a research society, CRA has closely followed the HPCC Act from its original inception and through its implementation as a program, and I am happy to have the opportunity to comment on both its past and its future.

I. Overview

Before commenting in detail on the HPCC Act, I would like to start my testimony with two general points.

• **First, we endorse strongly the findings of the National Research Council report referred to above and discussed by Dr. Sutherland.** In particular, we would emphasize the following points:

1. There has been mind-boggling progress in information technology, which pervades most aspects of our lives and of our economy.
2. The nation that leads in information technology enjoys enormous competitive advantages.
3. America owns this leadership today, thanks to a successful and complex interplay between government, academia, and industry in support of research. The track record is crystal clear.

4. The government's role is also crystal clear. Industry can afford to look ahead only a few years, but as a nation, we can and must invest for the long term. By and large, this fundamental basic research takes place in universities, with government support.

5. The HPCC program is a major success. In particular:

- The emphasis on "high-performance" is appropriate: cutting-edge information technology is a window on the future -- a "time machine." The high-performance technology of today will be the home, office, or schoolroom computer of 10 years from now.

- The emphasis on parallel computing is a success: although there is much more to be done, its viability is clear. Many important problems throughout science and engineering have been tackled, and nearly every vendor today has some sort of multiprocessor offering.

- The interagency coordination and cooperation is working exceedingly well. More than any other scientific field, computing research cuts across many agencies. There is no clear "lead" agency in computing, such as the role that the National Institutes of Health plays in biomedical research. HPCC represented a bold attempt to create, in a sense, a "virtual agency," a mechanism for coordinating the programs of diverse agencies, which serve diverse missions, but which all have some interest in advancing the state of computing and communication technology. From our perspective outside of government, this experiment appears to have worked exceedingly well, providing greater coherence and direction to the various programs with which we work.

- **Second, CRA believes that continued authorization of the HPCC Program--continued funding in these critical research areas and continued strengthening of the interagency**

process—is essential to the nation for the following reasons:

1. The HPCC program *is* the nation's research and education program in information technology. It is a coordinated multiagency initiative that supports nearly all of our nation's fundamental research and graduate education in information technology. HPCC is much more than the support of supercomputer centers, although the centers have grown into multidimensional institutions that make a wide variety of contributions to science and engineering. HPCC is much more than research on the highest-performance machines, although these systems are indeed "time machines" that offer an invaluable window into the future. Instead, HPCC is--and always has been--systems, software, networking, human resources, and research on information infrastructure technology and applications.
2. The interagency coordination fostered by the HPCC Act has been a model success in coordinating the efforts of multiple agencies with multiple approaches. This close coordination can be seen in several specific program initiatives, including the supercomputer centers, the gigabit testbeds, the digital library initiative, and the CIC strategic plan.
3. The HPCC program has proven to be appropriately flexible and adaptable. Fundamental research is inherently unpredictable. When Lewis and Clark were exploring America, they had a strategic goal, but their precise path of exploration was necessarily marked by false starts, backtracking, changes in direction, and searching for new paths. And so it has been with HPCC. The program has adapted and evolved in many ways, including:
 - increased focus on software.
 - increased focus on high speed digital communications.
 - the addition of the Information Infrastructure Technology and Applications element to the program structure.

4. A strategic plan for the future exists. The cooperation and coordination stimulated by the HPCC program has been extended to a strategic planning effort on the part of the National Science and Technology Council's Committee on Information and Communications, chaired by Dr. Anita K. Jones, Defense Director of Research and Engineering, and co-chaired by Dr. Paul Young, NSF's Assistant Director for Computer and Information Science and Engineering. The plan, *America in the Age of Information*, identified six Strategic Focus Areas "to focus fundamental information and communications research and to accelerate development in ways that are responsive to NSTC's overarching goals, agency mission goals, and our Nation's long term economic and defense needs." These Strategic Focus Areas are global-scale information infrastructure technologies, high performance / scalable systems, high-confidence systems, virtual environments, user-centered interfaces and tools, and human resources and education.

We think such multiagency planning and program implementation efforts are excellent, and we applaud them.

5. Finally, the role of universities is critical. The track record is clear. Federally supported university research has played a key role in essentially every aspect of modern information technology: timesharing, computer networking (the Internet), high-power workstations, computer graphics, database technology, Very Large Scale Integrated circuit design (VLSI), Reduced Instruction Set (RISC) processor architectures, Input/Output systems based on Redundant Arrays of Inexpensive Disks (RAID), parallel computing, and so on.

Universities look to the future. The HPCC program has been a huge success in allowing them to push the frontiers of their research further into the future. It is also important to emphasize that university research carried out under HPCC avoids picking "winners and losers." The purpose of publicly funded research in high-technology fields is to advance basic knowledge and create new opportunities that, in the medium and long term, industry exploits.

II. The HPCC Act

Now, let me speak about the particular aspects of the HPCC Act and offer some thoughts about future directions. The Act was organized around four principal sections, (1) systems, (2) software, (3) NREN, and (4) Basic Research and Human resources. Since that time, the program has also been expanded to a fifth component, Information Infrastructure Technology and Applications, which emphasizes research on leading-edge applications of information systems in areas of high potential impact, areas such as education, libraries, public health, and the like as well as the technological infrastructure to support them.

1. Systems: When the HPCC program was first being formulated in Congress and in the Administration, it was focused primarily on the largest high-end systems known as "supercomputers," but even at the time the HPCC Act was enacted, the focus of the program was broadening to a wider range of architectural goals. The emphasis was not so much on bigger, faster systems based on traditional architectures, but on exploring new, experimental architectures. In particular, researchers were exploring basic questions about the viability of the highly-parallel, scalable computer systems, and examining many different architectural concepts that seemed to have merit. Now we have a much better idea of the most promising architectural lines.

Still, significant challenges remain, particularly how to scale parallel systems to higher numbers of processors. Performance at the chip level is improving at 50% per year, meaning that the performance potential far outweighs our knowledge of how to assemble them together into productive systems and manage the flow of work through them.

2. Software: The focus at the time was on Grand Challenge computational science. In the

intervening years, progress has been made on many computationally demanding areas of basic scientific research.

Over the last four years, however, our vision has broadened substantially in two ways. In the first place, as we have progressed in our understanding of the design of scalable parallel architectures, there is an ever greater need to progress on our fundamental understanding of software and algorithms. Although important challenges remain on the hardware side, a proportionally greater emphasis is needed on software. If, as I suggested earlier, these advanced architectures will likely be the basis of everyday desktop systems of the next decade (and probably well beyond), research undertaken now on software and algorithms for these leading-edge systems will build the foundation for using them efficiently and effectively in the next century.

Although computational science “Grand Challenges” remain exciting and important to explore, we are now looking at a much wider range of “National Challenges,” applications that are crucial to the evolution of the nation’s information infrastructure.

3. National Research and Education Network (NREN): What a success story! Five years ago, the Internet was still pretty much an academic research network, used by researchers and students at universities. Even then, however, the base of users was broadening to undergraduate schools, to libraries, to K-12 education and to civic networks, so-called “freenets.” It was at the time far from clear that NREN would ever be much more than such a specialized system.

Now, NSF has nearly completed the process of spinning the Internet off to the private sector. It continues to grow at an explosive rate. Newspapers and magazines carry articles every day about the Internet and the World Wide Web. Commercial firms are fighting each other in the courts over network domain names. Packet-switching communications technology is an important component of the communications service and hardware industry.

NSF, in its concern for the health of U.S. science, needs to ensure that, as the Internet becomes commercialized, the needs of researchers and students for specialized advanced data-communication services are met. A major research responsibility also remains. As fast as researchers find ways to increase the speed of networks, both the growth of traffic and the demands of new applications find ways to consume resources. Thus, there remains an ongoing research agenda in extremely high-speed, extremely large-scale data networking, an agenda that should remain in the next generation Act. Such an agenda would include the following:

- **Scaling:** For all its growth, the Internet is still relatively small compared with an information infrastructure that would serve all of our society. Problems of scaling become even more complex as applications, such as the World Wide Web, are developed that gobble up increasing amounts of communications resources.
- **Quality of service:** The Internet was originally designed and built as a research network. To be fully usable in a commercial and public arena, its performance must be made more predictable and reliable.
- **Security:** If valuable assets or sensitive private information are to be transferred over the network, it must be made more secure, and ways to protect information in an insecure environment need to be developed.
- **Accounting:** Although this may seem to be a simple administrative matter, we have only begun to explore how to measure, manage, and account for the flow of information through extremely high-speed networks. Without new insights into how to accomplish this task with minimal overhead, such processes could more than double the cost of data transmissions, simply to allow for accounting.

- Technical/legal issues: Many policy and legal issues are being raised as the infrastructure moves more broadly into commercial use, issues such as intellectual property protection, privacy, access to government information, the First Amendment and censorship, and the like. Many of these issues have at least in part a technological component. Some careful research could help by identifying possible technological solutions or by clarifying the technological nature of the problem.

4. Basic Research and Human Resources: When the original bill was being considered, CRA pointed out repeatedly the need for a focus on basic research as the necessary underpinning for any HPCC program. So strongly did we feel about this that we insisted that basic research needed to be specifically identified in a fourth section of the bill, even though the other three sections arguably could be read to include fundamental investigations. Happily, Congress agreed with us.

We think that the need for strong support of basic research is unchanged. If anything, the focus is even more on the need for such fundamental work.

CRA has on its home page, pointers to a set of case studies describing how basic academic research resulted in significant application areas and economic benefits. I have attached four of them as an appendix to this testimony. We hope to continually add to and update this set of examples, which clearly makes the case for basic computing research.

5. Information Infrastructure Technology and Applications: As discussed above, this new area of interest has emerged as a fifth program element. I'll elaborate briefly on just two examples of National Challenges: educational technology and digital libraries.

Interest is growing in educational technology. I note that the Committee on Science held joint hearings with the Committee on Human Resources earlier this month to explore the potential of information technology to transform education. In those hearings, computer scientists such as Seymour Papert drew an ambitious and futuristic vision of how information technology could fundamentally transform learning. Whether Papert's vision is correct or not in the details, there is no question that information systems in the future will have the potential to play an enormous role in education. Nor do we doubt that a lot of fundamental computing research needs to be pursued before we can tap that potential.

Just one month ago, with NSF sponsorship, CRA held a two-day workshop on a basic computing research agenda for education. We called together about 100 computer scientists and education researchers to discuss the long-term needs. A preliminary report will be completed in a month or so, and we expect the final version, which will be published both in paper form and electronically on the Web, to be available early next year. We would be glad to brief the Science Committee and/or staff on these reports and our findings.

Similarly, in the area of digital libraries, we face a significant basic research agenda. For all the excitement that rightfully attended the evolution of the World Wide Web, information on the web has been likened to taking all the books in a large library and dumping them at random on the floor. We really don't know how to organize and search for information in such a massive distributed environment. We don't know how best to display it, how to use it, how to protect intellectual property rights for proprietary data while maintaining access to public information, and how to protect the privacy of users.

III. Recommendations

To summarize our recommendations on the next-generation HPCC Act:

1. We believe that a reauthorization of the High Performance Computing and Communications Act would be an important statement by Congress of the need to continue the nation's long-standing commitment to fundamental research in the computer field. It would serve to set general priorities within the computer field and provide a foundation for interagency coordination. The Computing Research Association would support such a bill.
2. The focus of the research program needs to be broad, but concentrated on fundamental research on the design and use of advanced leading-edge parallel, scalable computer systems, on extremely-high-speed data communications, and on the connections between the two. In particular, there is an important basic research agenda that underlies so-called "National Challenges," that will be a basic framework for the nation's infrastructure--education, library, health, government services, and so on.
3. The authorization needs to be flexibly drawn, allowing the program to adapt quickly as new research targets of opportunity appear.

Mr. Chairman, we understand that money is tight and that there are many claims on it. We also understand the enormous pressures the Congress finds itself in when trying to preserve research funding. But, we are also beginning to realize an enormous return on 50 years of investment in computing research (next year, we will celebrate the 50th anniversary of the invention of the stored program computer.) The information industry is still the fastest growing and now biggest sector of our economy.

We cannot let our investments in the basic research that underpins this field falter. If we fail to invest in research in information technology today, we will lose our leadership tomorrow, and once lost, it would be difficult, if not impossible, to recapture.

We commend your interest. Should the committee decide to draft a new bill, CRA looks forward to working with you in preparing a bill that will be an effective and worthy follow-on to the original.



Computing Research Association

Computing Research: Driving Information Technology and the Information Industry Forward

Information technology is central to our economy and to our society. It drives many of today's innovations and it offers enormous potential for further innovation in the coming decades. It also is the basis for an extremely successful \$500 billion industry that is critical to our nation's international competitiveness.

America still holds a commanding lead in this arena. This lead is the result of an extraordinary 50-year partnership among government, industry, and academia, in which federally-sponsored university research has played a critical role. The contributions of the Department of Defense Advanced Research Projects Agency and of the National Science Foundation are particularly notable for the way in which they have nourished the ideas and people that have let industry flourish.

But federal support for research in information technology is in jeopardy -- we are in danger of killing the golden goose. Some of the contributing factors are not specific to the field: the need to reduce deficit spending, an active debate concerning the appropriate role of federal funding in a wide range of areas, and concerns about the efficiency of university research. However, several misunderstandings specific to the information technology field also are playing a role:

- It is important that research in information technology not be caught in the middle of a misguided debate over "basic" versus "applied" research. As a recent study by the National Research Council shows, research in information technology occupies a complex middle ground, which can be characterized as "fundamental research in support of strategic directions." Such research advances our fundamental understanding of information, computing and communications, and at the same time creates new opportunities that industry can exploit in the medium and long term. However, because this research ultimately helps to renew existing companies and create new ones, its fundamental nature is sometimes missed, and support for such research is sometimes misconstrued as "industrial policy."
- The history of innovation in information technology needs to be better understood. Because the financial rewards in the information technology industry are so great, it is tempting to believe that industry must have been responsible for most of the fundamental innovations, and can be relied upon exclusively in the future. In fact, the contributions of federally-sponsored university research are pivotal. (Examples of these contributions are a key focus of this presentation.) Furthermore, international competitive pressures are driving companies to focus increasingly on near-term development with short-term payoffs in their own R&D, rather than on fundamental advances that lay the foundation for the future. As a result, America is increasingly reliant on universities for medium- and long-range R&D, for ideas, and for trained professionals.

There can be no doubt that it is necessary to reduce deficit spending sharply. There can be no doubt that it is appropriate to debate the role of the federal government in a broad spectrum of activities. And there can be no doubt that America's university system, for all of its remarkable accomplishments, is not above reproach or improvement. But there must also be no doubt about the role that federally-sponsored research in information technology has played, is playing, and must continue to play in our economy and in our society.

The purpose of this presentation is to document the importance of these issues, and to provide a factual

foundation for their discussion. We do this in two ways. First, we attach four brief essays. These essays are written by leading figures in the information industry -- people with decades of experience and perspective. They describe in an accessible way the history, impact, and promise of information technology research in four major areas (*please click to view these attachments*):

- [Database systems](#) (by James N. Gray, Senior Researcher, Microsoft Corporation)
- [Computer graphics](#) (by Edward R. McCracken, Chairman and Chief Executive Officer, Silicon Graphics, Inc.)
- [Reduced Instruction Set Computers \(RISC\)](#) (by William N. Joy, Co-Founder and Vice President for Research, Sun Microsystems, Inc.)
- [Computer communication networks](#) (by Vinton G. Cerf, Senior Vice President, Data Services Division, MCI Telecommunications Corporation)

Second, we elaborate on some of the issues raised above (*please click to view the text of these points*):

- [Information technology is central to our society -- economically and socially.](#)
- [America's leadership in information technology is the result of an extraordinarily complex and fruitful long-term partnership among government, industry, and academia.](#)
- [The government component of this partnership -- the highly effective public research program sponsored by ARPA and NSF -- cannot be pigeon-holed as "basic" or "applied" -- it is best characterized as "fundamental research in support of strategic directions."](#)
- [Enormous possibilities still lie ahead -- in fact, the real information revolution is yet to come.](#)
- [Continued federal investment is critical to realizing this promise -- to continuing our nation's momentum and to retaining our lead.](#)

Information technology and the information industry are driven forward in large part by the *ideas and people* that flow from the ARPA and NSF programs of fundamental research in support of strategic directions. Today, the High Performance Computing and Communications Initiative (HPCCI) is the multiagency cooperative effort under which this work moves forward. Led by the [National Coordination Office](#), HPCCI's accomplishments are many-fold (see the [FY 1996 "Blue Book"](#) for a summary). And the need and the potential for research in information technology has never been greater. Programs like HPCCI -- particularly with its increasing emphasis on enabling broad-based future uses of information infrastructure (see [America in the Age of Information](#), the strategic plan recently published by the National Science and Technology Council's Committee on Information and Communications) -- are preparing the nation for the 21st century. Such programs require -- and have earned -- the strong support of Congress and of the American people.

Please click to view a list of related material.

Copyright 1995 by Edward D. Lazowska and the Computing Research Association. Contributions to this document by John V. Guttag and Fred W. Weingarten are appreciated. Portions adapted with permission from "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure", copyright 1995 by the National Academy of Sciences. Courtesy of the National Academy Press, Washington DC.

<http://cra.org/research.impact>
(better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
lazowska@cs.washington.edu

Database Systems: A Textbook Case of Research Paying Off



James N. Gray
Senior Researcher
Microsoft Corporation

Industry Profile

The database industry generated about \$7 billion in revenue in 1994 and is growing at 35% per year. Among software industries, it is second only to operating system software. All of the leading corporations in this industry are US-based: IBM, Oracle, Sybase, Informix, Computer Associates, and Microsoft. In addition, there are two large specialty vendors, both also US-based: Tandem, selling over \$1 billion per year of fault-tolerant transaction processing systems, and AT&T-Teradata, selling about \$500 million per year of data mining systems.

In addition to these well-established companies, there is a vibrant group of small companies specializing in application-specific databases -- text retrieval, spatial and geographical data, scientific data, image data, and so on. An emerging group of companies offer object-oriented databases. Desktop databases are another important market focused on extreme ease-of-use, small size, and disconnected operation.

A relatively modest federal research investment, complemented by an also-modest industrial research investment, has led directly to our nation's dominance of this key industry.

Historical Perspective

Companies began automating their back-office bookkeeping in the 1960s. COBOL and its record-oriented file model were the work-horses of this effort. Typically, a batch of transactions was applied to the old-tape-master, producing a new-tape-master and printout for the next business day.

During this era, there was considerable experimentation with systems to manage an on-line database that could capture transactions as they happened, rather than in daily batches. At first these systems were ad hoc, but late in the decade "network" and "hierarchical" database products emerged. A network data model standard (DBTG) was defined, which formed the basis for most commercial systems during the 1970s. Indeed, in 1980 DBTG-based Cullinet was the leading software company.

However, there were some problems with DBTG. DBTG used a low-level, record-at-a-time procedural

language. The programmer had to navigate through the database, following pointers from record to record. If the database was redesigned, as they often are over a decade, then all the old programs had to be rewritten.

The "relational" data model, enunciated by Ted Codd in a landmark 1970 article, was a major advance over DBTG. The relational model unified data and metadata so that there was only one form of data representation. It defined a non-procedural data access language based on algebra or logic. It was easier for end-users to visualize and understand than the pointers-and-records-based DBTG model. Programs could be written in terms of the "abstract model" of the data, rather than the actual database design; thus, programs were insensitive to changes in the database design.

The research community (both industry and university) embraced the relational data model and extended it during the 1970s. Most significantly, researchers showed that a high-level relational database query language could give performance comparable to the best record-oriented database systems. This research produced a generation of systems and people that formed the basis for IBM's DB2, Ingres, Sybase, Oracle, Informix and others. The SQL relational database language was standardized between 1982 and 1986. By 1990, virtually all database systems provided an SQL interface (including network, hierarchical and object-oriented database systems, in addition to relational systems).

Meanwhile the database research agenda moved on to geographically distributed databases and to parallel data access. Theoretical work on distributed databases led to prototypes which in turn led to products. Today, all the major database systems offer the ability to distribute and replicate data among nodes of a computer network.

Research of the 1980s also showed how to execute each of the relational data operators in parallel -- giving hundred-fold and thousand-fold speedups. The results of this research are now beginning to appear in the products of several major database companies.

Three Case Studies

The government has funded a number of database research efforts from 1970 to the present. Projects at UCLA gave rise to Teradata and produced many excellent students. Projects at CCA (SDD-1, Daplex, Multibase, and HIPAC) pioneered distributed database technology and object-oriented database technology. Projects at Stanford created deductive database technology, data integration technology, and query optimization technology. Work at CMU gave rise to general transaction models and ultimately to the Transarc Corporation. There have been many other successes from AT&T, the University of Texas at Austin, Brown, Harvard, Maryland, Michigan, MIT, Princeton, and Toronto, among others. It is not possible to enumerate all the contributions here, but we shall highlight three representative research projects that had major impact on the industry.

Ingres

Project Ingres started at UC Berkeley in 1972. Inspired by Codd's work on the relational database model, several faculty members (Stonebraker, Rowe, Wong, and others) started a project to design and build a relational database system. Incidental to this work, they invented a query language (QUEL), relational optimization techniques, a language binding technique, and interesting storage strategies. They also pioneered work on distributed databases.

The Ingres academic system formed the basis for the Ingres product now owned by Computer Associates. Students trained on Ingres went on to start or staff all the major database companies (AT&T, Britton Lee, HP, Informix, IBM, Oracle, Tandem, Sybase). The Ingres project went on to investigate distributed databases, database inference, active databases, and extensible databases. It was rechristened Postgres, which is now the basis of the digital library and scientific database efforts within the University of California system. Recently, Postgres spun off to become the basis for a new object-relational system from the startup Illustra Information Technologies.

System R

Codd's ideas were inspired by seeing the problems IBM and its customers were having with the DBTG network data model and with IBM's product based on this model (IMS). Codd's relational model was at first

very controversial; people thought that the model was too simplistic and that it could never give good performance. IBM Research management took a gamble and chartered a 10-person effort to prototype a relational system based on Codd's ideas. This group produced a prototype, System R, that eventually grew into the DB2 product series. Along the way, the IBM team pioneered ideas in query optimization, data independence (views), transactions (logging and locking), and security (the grant-revoke model). In addition, the SQL query language from System R was the basis for the standard that emerged.

The System R group went on to investigate distributed databases (project R*) and object-oriented extensible databases (project Starburst). These research projects have pioneered new ideas and algorithms. The results appear in IBM's database products and in those of other vendors.

Gamma

During the 1970s there was great enthusiasm for database machines -- special-purpose computers that would be much faster than general-purpose systems running conventional databases. The problem was that general purpose systems were improving at 50% per year, so it was difficult for customized systems to compete with them. By 1980, most researchers recognized the futility of special-purpose approaches, and the database machine community switched to research on using arrays of general purpose processors and disks to process data in parallel. The University of Wisconsin was home to the major proponents of this idea in the US. Funded by the government and industry, they built a parallel database machine called Gamma. That system produced ideas and a generation of students who went on to staff all the database vendors. Today the parallel systems from IBM, Tandem, Oracle, Informix, Sybase, and AT&T all have a direct lineage from the Wisconsin research on parallel database systems. The use of parallel database systems for data mining is the fastest-growing component of the database server industry.

The Gamma project evolved into the Exodus project at Wisconsin (focusing on an extensible object oriented database). Exodus has now evolved to the Paradise system which combines object-oriented and parallel database techniques to represent, store, and quickly process huge earth-observing satellite databases.

The Future

Database systems continue to be a key aspect of Computer Science & Engineering today. Representing knowledge within a computer is one of the central challenges of the field. Database research has focused primarily on this fundamental issue. Many universities have faculty investigating these problems and offer courses that teach the concepts developed by this research program.

There continues to be active and valuable research on representing and indexing data, adding inference to data search, compiling queries more efficiently, executing queries in parallel, integrating data from heterogeneous data sources, analyzing performance, and extending the transaction model to handle long transactions and workflow (transactions that involve human as well as computer steps). The availability of very-large-scale (tertiary) storage devices has prompted the study of models for queries on very slow devices.

In addition, there is great interest in unifying object-oriented concepts with the relational model. New datatypes (image, document, drawing) are best viewed as the methods that implement them rather than the bytes that represent them. By adding procedures to the database system, one gets active databases, data inference, and data encapsulation. This object-oriented approach is an area of active research and ferment both in academe and in industry.

A very modest research investment produced American market dominance in a \$7 billion industry -- creating the ideas for the current generation of products, and training the people who built those products. Continuing research is creating the ideas and training the people for the next product generation.

The Author

Dr. James N. Gray is one of the world's leading experts on database and transaction processing computer systems. Over the past three decades he has worked for IBM, Tandem, and Digital Equipment Corporation on systems including System R, SQL/DS, DB2, IMS-Fast Path, Encompass, NonStopSQL, Pathway, TMF,

Rdb, DBI, and ACMS -- systems that have defined the progress of the field.

Dr. Gray holds doctorates from U.C. Berkeley and the University of Stuttgart. He is a Fellow of the ACM, a member of the National Research Council's Computer Science and Telecommunications Board, Editor-in-Chief of the *VLDB Journal*, Editor of the Morgan Kaufmann series on Data Management, editor of the *Performance Handbook for Database and Transaction Processing Systems*, co-author of *Transaction Processing Concepts and Techniques*, and serves on Objectivity's Technical Advisory Board.

After six months as McKay Fellow in U.C. Berkeley's Computer Science Division, he has just joined Microsoft to establish a San Francisco Bay Area laboratory focusing on making Microsoft data servers more scalable, manageable, and fault tolerant.

Back to [Computing Research: Driving Information Technology and the Information Industry Forward](#)

Copyright 1995 by James N. Gray and the Computing Research Association. Contributions to this document by Philip A. Bernstein, David DeWitt, Michael Stonebraker, and Jeffrey D. Ullman are appreciated.

<http://cra.org/research.impact>
(better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
lazowska@cs.washington.edu

Computer Graphics: Ideas and People from America's Universities Fuel a Multi-Billion-Dollar Industry



Edward R. McCracken
Chairman and Chief Executive Officer
Silicon Graphics, Inc.

Overview

Advances in computer graphics have transformed how we use computers. Computer graphics has given us the "mouse" input device, "what-you-see-is-what-you-get" document preparation systems, the computer-aided design systems used to create the Boeing 777, the ability to visualize molecular dynamics and other scientific phenomena, the animation used in educational software and the advertising and entertainment industries, and virtual reality systems whose applications range from architectural prototyping to surgical training to entertainment. Today, every user of a computer benefits from computer graphics, even in applications such as word processors, spreadsheets, databases, and project planners. Because of user-friendly graphical user interfaces, pre-schoolers now routinely use computers, a revolution undreamt of even a few years ago.

While everyone is familiar with the mouse, multiple "windows" on computer screens, and stunningly realistic images of everything from animated logos in television advertisements to NASA animations of spacecraft flying past Saturn, few people realize that these innovations were spawned by federally sponsored university research. Without far-sighted support from agencies such as the Department of Defense Advanced Research Projects Agency and the National Science Foundation, computer graphics and the multi-billion-dollar industry it makes possible would have developed much more slowly, and perhaps not predominantly in the U.S.

Historical Perspective

From its beginnings in the late 1960s, when a few ARPA- and NSF-sponsored research laboratories were working on relatively obscure graphics-related projects, the computer graphics community has grown to more than a hundred thousand software and hardware engineers and application developers. Early thrusts included graphics support software and rendering algorithms, graphics hardware architectures, graphical user interfaces, and hypermedia. The availability of graphical tools and systems has vastly influenced developments in computer-aided design and manufacturing, including the automotive and aerospace industries, molecular modeling and drug design, medical imaging, architectural design, and the entertainment industry. Today, many scientific and engineering disciplines that were once distinct from computer graphics are inextricably

interwoven with it. In these disciplines, visualization is no longer an optional tool but a critical enabling technology.

By far the largest segments of today's computing industry are the personal computer and workstation markets. All major vendors (IBM, HP, Sun, DEC, SGI, Intel, Apple) participate in these segments, which total roughly \$50 billion and \$15 billion, respectively. And note that these dollar volumes only represent the hardware and system-software portions of these markets; system software is a small share of the total software market, which is dominated by applications, almost all of which make use of computer graphics.

Case Studies

The nation should look back with great pride on its research investments in all major areas of computer graphics. Six of the most significant of these areas are discussed below.

User Interfaces

The power of computers cannot be harnessed without a way to access and control that power; the interface between the user and the machine can determine the success or failure of both hardware and software. Apple's graphical desktop interface for the Macintosh computer (and the Microsoft Windows equivalent for PCs), and more recently the introduction of Mosaic, a graphical browser for the Internet, are excellent examples of how applications of computer graphics research can create new markets and broaden old ones.

Both of these applications of graphics technology -- the desktop metaphor and the Mosaic browser -- had their origins in federally sponsored efforts. ARPA-sponsored research at the University of Utah was built on by Alan Kay in the 1970s at Xerox PARC to create the Smalltalk programming environment on the pioneering Alto bitmapped graphics workstation. This environment and PARC's Bravo document editor stimulated the development of the Apple Macintosh (1984) and bitmapped graphically based windowing systems and graphical user interfaces. NSF funding supported the development of the Mosaic browser at the National Center for Supercomputing Applications. Now, K-12 students throughout America use computers as information access devices -- they treat the Internet as a digital library and truly have information "at their fingertips." Mosaic has not only exponentially increased the number of Internet users, but is also spawning many new companies and enterprises and is currently arousing intense corporate interest.

Computer Graphics Hardware

The hardware used in interactive computer graphics has its genesis in federally sponsored university research. The industry leader in rendering hardware is Silicon Graphics, Inc., founded by Jim Clark. Clark received his Ph.D. from the University of Utah where he and his advisor, Ivan Sutherland, pursued a federally funded program of research in 3D graphics hardware. Joining the faculty at Stanford, Clark received support from the ARPA VLSI Program for his Geometry Engine project, whose goal was to harness modern custom integrated-circuit technology to create cost-effective high-performance graphics systems. It was this Geometry Engine that formed the basis of SGI.

In 1968, Douglas Engelbart of Stanford Research Institute demonstrated his hypertext system, NLS, which was funded by ARPA. Among other things, this system included the first mouse -- now a standard fixture of computer systems everywhere. (One of Microsoft Corporation's highest-revenue products is the Microsoft Mouse!)

Hypertext/Hypermedia

Hypertext and hypermedia have their roots in Vannevar Bush's famous 1945 *Atlantic Monthly* article, "As We May Think." Bush described how documents might be interlinked in the fashion of human associative memory. These ideas inspired Doug Engelbart at SRI (funded by ARPA) and Andries van Dam of Brown University (funded by NSF) to develop the first hypertext systems in the 1960s. These systems were the forerunners of today's word-processing programs, including simple what-you-see-is-what-you-get capabilities that were further refined in the Xerox Bravo editor. The ideas and concepts were fundamental to such developments as Apple's popular Hypercard and NCSA Mosaic.

Rendering

High-quality rendering has caught the public's eye and is having a vast impact on the entertainment and advertising industries. From *Jurassic Park* to simulator rides at Disney World and dancing soda cans in TV commercials, the world has been seduced by computer animation, special effects, and photorealistic imagery of virtual environments. How are these pictures created and where did the techniques for creating them come from?

ARPA and NSF deserve the lion's share of the credit. Before there was a market and a demand, they supported research activities at the University of Utah, North Carolina State, Ohio State, Caltech, and Cornell. For example, Gouraud shading, which allowed substantially more realistic images than the previous wire-frame images, was developed in 1970 by Henri Gouraud, a graduate student at the University of Utah. Coupled with work at the New York Institute of Technology, these results provided the foundation for the sophisticated software commonly available today. The industry still uses the basic algorithms developed at Utah in the 1970s for simple lighting calculations, both in software and increasingly in commodity hardware. More sophisticated rendering packages exploiting university-developed algorithms, such as Pixar's Renderman, are used by the film and animation industries, as well as in flight simulators and automotive design.

Graphics Software Systems

The combination of the above advances led to many commercial graphics software systems. Rather than develop new systems for each application area -- be it advertising, animation, molecular modeling, or scientific visualization -- it began to make sense to utilize general-purpose systems. Offerings from Wavefront, AVS, SGI (Iris Explorer), and IBM (Data Explorer) were developed with strong influence by former Cornell and Brown students, educated in NSF-sponsored graphics laboratories. PostScript, the de facto standard in page-description languages for laser printers, was developed by Adobe, founded by Utah Ph.D. John Warnock. That graduates of federally sponsored university graphics research laboratories move on to lead industrial projects demonstrates the most effective means of technology transfer between universities and industry.

Virtual Reality

The popular idea of virtual reality saw its first implementation in Ivan Sutherland's ground-breaking work at Harvard in 1968. With funding from both commercial and government sources, including ONR, Bell Laboratories, the US Air Force and the CIA, Sutherland's work included the first head-mounted display as well as stereo and see-through displays, head tracking, and a hand-held 3D cursor. Such devices have now become widespread and are used in areas as diverse as video game systems, rapid prototyping for industrial design and architecture, and scientific visualization. Boeing's new 777 airplane was designed electronically throughout, including CAD/CAM 3D models, windtunnel simulation, and virtual-reality-based accessibility studies. This digital design enabled Boeing to avoid \$100 million mockups, and the plane came together with far fewer changes and far greater accuracy than any previous design, enabling Boeing to maintain its competitive edge. Sutherland's early VR work also had influence on flight simulators, and his company, Evans and Sutherland, Inc., pioneered the visual simulation market, now a major business for many companies.

The Future

One could continue with many more examples, but the message is clear: federal sponsorship of university research in computer graphics stimulated a major segment of the computing industry, allowing the United States to establish and maintain a competitive edge.

We now are facing the next hurdle. As we move into the next century with ubiquitous computing, parallel processing, almost infinite bandwidth, a vastly broader community of users, telepresence, and collaborative computing, things will change drastically. We need to investigate and develop new parallel architectures, new pixel-based display architectures, better and faster rendering algorithms, easier-to-use 3D user interfaces, and physically based models and simulations. This research depends not only on computer science issues but also on physics, optics, thermodynamics, digital sampling theory, perception, human factors, graphical design, and other areas too numerous to list. In addition, the results are intertwined with application needs in manufacturing, medical imaging, molecular modeling, scientific visualization, pilot training, aircraft and automotive design, and education.

America's leadership in computer graphics, and in information technology as a whole, is the result of a remarkable long-term partnership among government, industry, and academia, in which federally sponsored university research plays a critical role. Now, more than ever, we must invest to maintain this leadership.

The Author

Edward R. McCracken is Chairman and Chief Executive Officer of Silicon Graphics, Inc., the leading manufacturer of high-performance visual computing systems. The company delivers interactive three-dimensional graphics, digital media, and multiprocessing supercomputing technologies to technical, scientific, and creative professionals, and is actively involved in bringing these same technologies to the home market through its relationships with AT&T, Nintendo, NTT, and Time Warner. Its subsidiary, MIPS Technologies, Inc., designs and licenses RISC processor technology for the computer systems and consumer electronics markets. The company's products have been used to create special effects in a number of recent movies, including *Jurassic Park* and *Forrest Gump*.

SGI has offices worldwide and headquarters in Mountain View, California. In the 12 months ending in March 1994, SGI reported net revenues of \$1.9 billion. The 13-year-old company employs 5,250 people.

McCracken has an MBA from Stanford University and a BSEE from Iowa State University. He serves on the Board of Directors of Tularik, Inc., a privately held biotechnology company. He also serves as co-chairman of Joint Venture: Silicon Valley Network, Inc., and co-chairman of the United States Advisory Council on the National Information Infrastructure.

Back to Computing Research: Driving Information Technology and the Information Industry Forward

Copyright 1995 by Edward R. McCracken and the Computing Research Association. Contributions to this document by Forest Baskett, Henry Fuchs, Donald P. Greenberg, and Andries van Dam are appreciated.

<http://cra.org/research.impact>
(better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
lazowska@cs.washington.edu

Reduced Instruction Set Computers (RISC): Academic/Industrial Interplay Drives Computer Performance Forward



William N. Joy
Co-Founder, and Vice President for Research
Sun Microsystems, Inc.

The Microprocessor Revolution

The engine of the computer revolution is the microprocessor. It has led to new inventions, such as FAX machines and personal computers, as well as adding intelligence to existing devices, such as wristwatches and automobiles. Moreover, its performance has improved by a factor of 25,000 since its birth in 1971! Has any other invention, so useful already at birth, undergone a similar improvement?

For the first 15 years of its life, the microprocessor improved its performance by an impressive 35% per year. As a result of research at universities and industrial laboratories, though, this rate increased in 1987 to about 55% per year -- a doubling of performance every 18 months! Had the prior rate of 35% been maintained, computers today would be only one third the speed we actually enjoy. Stated differently, it's as if we had been granted a wish to use the computers of the year 2000 today.

This increase coincided with the introduction of Reduced Instruction Set Computers (RISC). The *instruction set* is the hardware "language" in which the software tells the processor what to do. Surprisingly, reducing the size of the instruction set -- eliminating certain instructions based upon a careful quantitative analysis, and requiring these seldom-used instructions to be emulated in software -- can lead to higher performance, for several reasons:

- The vacated area of the chip can be used in ways that accelerate the performance of more commonly used instructions, more than compensating for the inevitably degraded performance of the seldom-used instructions.
- It becomes easier to optimize the design.
- It allows microprocessors to use techniques hitherto restricted to the largest computers.
- It simplifies translation from the high-level language in which people program into the instruction set that the hardware understands, resulting in a more efficient program.

RISC was also heralded a more quantitative approach to computer architecture, whereby careful experiments preceded the hardware design and sensible performance metrics were used to judge success. Previously,

intuition guided some computer design projects, leading to disappointing results. As the quantitative approach replaced the less well-defined prior approach, people in all divisions of a computer company were trained in this new school of design; hence engineering, manufacturing, marketing, sales, and management acted in concert, further accelerating the rate of change.

The Genealogy of RISC

The roots of RISC lie in three research projects: the IBM 801, the Berkeley RISC processor, and the Stanford MIPS processor. These architectures attracted enormous interest because of claims of a performance advantage of anywhere from two to five times over contemporary machines using traditional architectures.

Begun in the late 1970s, the IBM project was the first to start but was the last to become public. The IBM machine was designed as a minicomputer made from hundreds of chips, while the university projects were both microprocessors. John Cocke is considered to be the father of the 801 design. In recognition of his contribution he received both the Turing award, the highest award in computer science and engineering, and the Presidential Medal of Technology.

In 1980, David A. Patterson and his colleagues at the University of California at Berkeley, sponsored by the Department of Defense Advanced Research Projects Agency, began the project that was to give this approach its name. They built two machines, called RISC-I and RISC-II. Because the IBM project was not widely known or discussed, the role played by the Berkeley group in promoting the RISC approach was critical to the acceptance of the technology.

In 1981, John L. Hennessy and his colleagues at Stanford published a description of the Stanford MIPS machine, also developed under ARPA sponsorship. Both university projects were interested in designing a simple machine that could be built as a microchip within the university environment. All three early RISC machines had similar "reduced" languages.

Importantly, the Berkeley and Stanford projects fit within the broad mosaic of the overall ARPA VLSI Program, a highly ambitious program which envisioned that integrated circuit technology could be made available to system designers -- people with an overall view of the objectives and constraints of an entire hardware/software system -- and that this would have tremendous impact. The ARPA VLSI Program developed the concept of the multichip wafer, which allowed multiple integrated circuit designs to share a single silicon fabrication run, dramatically reducing costs. It conceived of the MOSIS fabrication service, which created a multichip wafer from designs submitted electronically from multiple sites, allowing university system designers access to state-of-the-art silicon fabrication. It sponsored extensive developments in computer-aided design tools. It launched, in addition to the RISC revolution, a comparable revolution in cost-effective high-performance computer graphics through the Geometry Engine and Pixel Planes projects, now the basis of Silicon Graphics, Inc., and Irix and Division. And it produced *people*.

Commercialization

In retrospect, it was important for the success of RISC for there to have been three research projects. The RISC ideas were highly controversial at first -- disbelievers were plentiful. In addition to exploring somewhat different avenues and making complementary discoveries, the three projects reinforced each others' results, and led to more voices advocating RISC in what was then a decidedly RISC-averse world.

In 1986 the computer industry began to announce commercial processors based on the technology explored by the three RISC research projects. Hennessy founded a company with the same name as his research project -- MIPS. Hewlett-Packard converted their existing minicomputer line to RISC architectures. IBM never directly turned the 801 into a product. Instead, the ideas were adopted for a new, low-end architecture that was incorporated in the IBM RT-PC. Alas, this machine was a commercial failure, but subsequent RISC processors in which IBM has been involved (e.g., the Apple/IBM/Motorola PowerPC) have been highly successful.

In 1987 Sun Microsystems began delivering machines based on the SPARC architecture, a derivative of the Berkeley RISC-II machine. In the view of many, it was Sun's success with RISC-based workstations that convinced the remaining skeptics that RISC was significant commercially. In particular, RISC advocates used Sun's success to get RISC restarted at IBM. IBM announced a new RISC architecture in 1990, as did DEC in

1993. Today, RISC is the foundation of a \$15 billion industry.

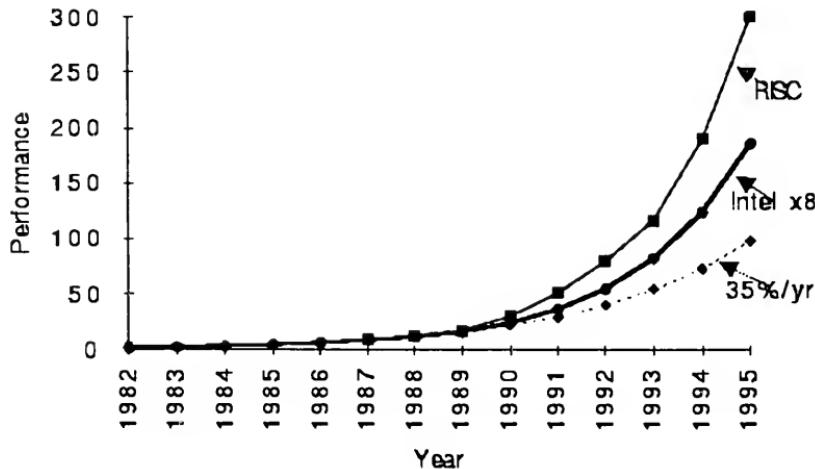
Intel's microprocessors are used in the popular IBM PC, and hence are the most widely used microprocessors, but they predate RISC. RISC microprocessors have been the standard-bearers of performance, so Intel has embraced ideas from RISC and followed the quantitative approach. Thus both the ideas and the competition from RISC has benefited all computer users, since RISC has raised the performance target for the entire industry. With the announcement that Hewlett-Packard and Intel will move to a common instruction set in 1997, the end of the non-RISC architectures draws near.

The Future

Today's microprocessors are almost 25,000 times faster than their ancestors. And microprocessor-based computer systems now cost only 1/40th as much as their ancestors, when inflation is considered. The result: an overall cost-performance improvement of 1,000,000! This extraordinary advance is why computing plays such a large role in today's world. Had the research at universities and industrial laboratories not occurred -- had the complex interplay between government, industry, and academia not been so successful -- a comparable advance would still be years away.

Microprocessor performance can continue to double every 18 months beyond the turn of the century. This rate can be sustained by continued research innovation. Significant new ideas will be needed in the next decade to continue the pace; such ideas are being developed by research groups today.

Looking further into the future, if this rate continues for the next 25 years, performance will improve by as much as it has in the last 50 years. The implications of such a breathtaking advance truly are limited only by our imaginations.



Microprocessor performance over time, relative to the CDC 6600, an early supercomputer

The Author

Dr. William N. Joy is a founder and currently Vice President for Research at Sun Microsystems, a leading manufacturer of powerful computer workstations. He is currently studying architectures for human-computer interaction, involving new kinds of interfaces, new system and application software architectures, and new ways of storing information to make information systems easier to use.

In the 1970s Joy was the principal designer of the U.C. Berkeley version of the UNIX operating system, whose networking protocols and implementations helped spawn the Internet. (Berkeley UNIX was an ARPA-sponsored effort.) In the 1980s he spearheaded Sun's evangelism of the "open systems" model of computing, which allows different groups to contribute to a computer system by making the specifications of its components freely available. Other technical contributions include co-design of the SPARC microprocessor architecture and several of its implementations, and design of the Network File System (NFS), an industry standard way for computers to share files.

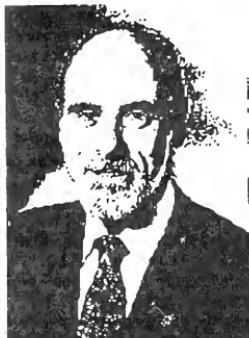
Joy is the recipient of the Grace Murray Hopper Award from the Association for Computing Machinery, and of the Lifetime Achievement Award from the USENIX Association. He has 11 patents issued or pending.

[Back to Computing Research: Driving Information Technology and the Information Industry Forward](#)

Copyright 1995 by William N. Joy and the Computing Research Association. Contributions to this document by John L. Hennessy and David A. Patterson are appreciated.

<http://cra.org/research.impact>
(better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
lazowska@cs.washington.edu

Computer Networking: Global Infrastructure for the 21st Century



Vinton G. Cerf
Senior Vice President, Data Services Division
MCI Telecommunications Corporation

The Internet Phenomenon

The Internet has gone from near-invisibility to near-ubiquity in little more than a year. (Near-ubiquity may be a bit of an over-statement, but Valvoline's promotion of its Internet World Wide Web home page in its Indianapolis 500 television advertisements speaks worlds about the phenomenon that we are witnessing!) In fact, though, today's multi-billion dollar industry in Internet hardware and software is the direct descendant of strategically-motivated fundamental research begun in the 1960s with federal sponsorship. A fertile mixture of high-risk ideas, stable research funding, visionary leadership, extraordinary grass-roots cooperation, and vigorous entrepreneurship has led to an emerging Global Information Infrastructure unlike anything that has ever existed.

Although not easy to estimate with accuracy, the 1994 data communications market approached roughly \$15 billion/year if one includes private line data services (\$9 billion/year), local area network and bridge/router equipment (\$3 billion/year), wide area network services (\$1 billion/year), electronic messaging and online services (\$1 billion/year), and proprietary networking software and hardware (\$1 billion/year). Some of these markets show annual growth rates in the 35-50% range, and the Internet itself has doubled in size each year since 1988.

The Internet now encompasses an estimated 50,000 networks worldwide, about half of which are in the United States. There are over 5 million computers permanently attached to the Internet, plus at least that many portable and desktop systems which are only intermittently online. (There were only 4 computers on the ARPANET in 1969, and only 200 on the Internet in 1983!) Traffic rates measured in the recently "retired" NSFNET backbone approached 20 trillion bytes per month and were growing at a 100% annual rate.

What triggered this phenomenon? What sustains it? How is its evolution managed? The answers to these questions have their roots in ARPA-sponsored research in the 1960s into a then-risky new approach to data communication: *packet switching*. The U.S. government has played a critical role in the evolution and application of advanced computer networking technology and deserves credit for stimulating wide-ranging exploration and experimentation over the course of several decades.

Evolutionary Stages

Packet Switching

Today's computer communication networks are based on a technology called *packet switching*. This technology, which arose from ARPA-sponsored research in the 1960s, is fundamentally different from the technology that was then employed by the telephone system (which was based on "circuit switching") or by the military messaging system (which was based on "message switching").

In a packet switching system, data to be communicated is broken into small chunks that are labeled to show where they come from and where they are to go, rather like postcards in the postal system. Like postcards, packets have a maximum length and are not necessarily reliable. Packets are forwarded from one computer to another until they arrive at their destination. If any are lost, they are re-sent by the originator. The recipient acknowledges receipt of packets to eliminate unnecessary re-transmissions.

The earliest packet switching research was sponsored by the Information Processing Techniques Office of the Department of Defense Advanced Research Projects Agency, which acted as a visionary force shaping the evolution of computer networking as a tool for coherent harnessing of far-flung computing resources. The first experiments were conducted around 1966. Shortly thereafter, similar work began at the National Physical Laboratory in the UK. In 1968 ARPA developed and released a Request for Quotation for a communication system based on a set of small, interconnected computers it called "Interface Message Processors" or "IMPs." The competition was won by Bolt Beranek and Newman (BBN), a research firm in Cambridge, MA, and by September 1969 BBN had developed and delivered the first IMP to the Network Measurement Center located at UCLA. The "ARPANET" was to touch off an explosion of networking research that continues to the present.

Apart from exercising leadership by issuing its RFQ for a system that many thought was simply not feasible (AT&T was particularly pessimistic), ARPA also set a crucial tone by making the research entirely unclassified and by engaging some of the most creative members of the computer science community who tackled this communication problem without the benefit of the experience (and hence bias) of traditional telephony groups. Even within the computer science community, though, the technical approach was not uniformly well-received, and it is to ARPA's credit that it persevered despite much advice to the contrary.

ARPANET

The ARPANET grew from four nodes in 1969 to roughly one hundred by 1975. In the course of this growth, a crucial public demonstration was held during the first International Conference on Computer Communication in October 1972. Many skeptics were converted by witnessing the responsiveness and robustness of the system. Out of that pivotal meeting came an International Network Working Group (INWG) composed of researchers who had begun to explore packet switching concepts in earnest. Several INWG participants went on to develop an international standard for packet communication known as X.25, and to lead the development of commercial packet switching in the U.S., Canada, France, and the UK, specifically for systems such as Telenet, Datapac, Experimental Packet Switching System, Transpac, and Reseau Communication par Paquet (RCP).

By mid-1975, ARPA had concluded that the ARPANET was stable and should be turned over to a separate agency for operational management. Responsibility was therefore transferred to the Defense Communications Agency (now known as the Defense Information Systems Agency).

New Packet Technologies

ARPANET was a single terrestrial network. Having seen that ARPANET was not only feasible but powerfully useful, ARPA began a series of research programs intended to extend the utility of packet switching to ships at sea and ground mobile units through the use of synchronous satellites (SATNET) and ground mobile packet radio (PRNET). These programs were begun in 1973, as was a prophetic effort known as "Internetting" which was intended to solve the problem of linking different kinds of packet networks together without requiring the users or their computers to know much about how packets moved from one network to another.

Also in the early 1970s, ARPA provided follow-on funding for a research project originated in the late 1960s

by the Air Force Office of Scientific Research to explore the use of radio for a packet switched network. This effort, at the University of Hawaii, led to new mobile packet radio ideas and also to the design of the now-famous Ethernet. The Ethernet concept arose when a researcher from Xerox PARC spent a sabbatical period at the University of Hawaii and had the insight that the random access radio system could be operated on a coaxial cable, but at data rates thousands of times faster than could then be supported over the air. Ethernet has become a cornerstone of the multi-billion dollar local area network industry.

These efforts came together in 1977 when a four-network demonstration was conducted linking ARPANET, SATNET, Ethernet and the PRNET. The satellite effort, in particular, drew international involvement from participants in the UK, Norway, and later Italy and Germany.

The Internet Protocols

A third ARPA effort of the early 1970s involved research at Stanford to design a new set of computer communication protocols that would allow multiple packet networks to be interconnected in a flexible and dynamic way. In defense settings, circumstances often prevented detailed planning for communication system deployment, and a dynamic, packet-oriented, multiple-network design provided the basis for a highly robust and flexible network to support command-and-control applications.

The first phase of this work culminated in a demonstration in July 1977, the success of which led to a sustained effort to implement robust versions of the basic Internet protocols (called TCP/IP for the two main protocols: Transmission Control Protocol and Internet Protocol). The roles of ARPA and the Defense Communications Agency were critical both in supplying sustained funding for implementing the protocols on various computers and operating systems and for the persistent and determined application of the new protocols to real needs.

By 1980, sufficient experience had been gained that the design of the protocols could be frozen and a serious effort mounted to require all computers on the ARPANET to adopt TCP/IP. This effort culminated in a switch to the new protocols in January 1983. ARPANET had graduated to production use, but it was still an evolving experimental testbed under the leadership of ARPA and DCA.

ARPANET -> NSFNET -> Internet

As ARPA and DCA were preparing to convert the organizations they supported to TCP/IP, the National Science Foundation started an effort called CSNET (for Computer Science Network) to interconnect the nation's computer science departments, many of which did not have access to ARPANET. CSNET adopted TCP/IP, but developed a dial-up "Phone-mail" capability for electronic mail exchange among computers that were not on ARPANET, and pioneered the use of TCP/IP over the X.25 protocol standard that emerged from commercial packet switching efforts. Thus, the beginning of the 1980s marked the expansion of U.S. government agency interest in networking, and by the mid-1980s the Department of Energy and NASA also had become involved.

NSF's interest in high-bandwidth attachment was ignited in 1986 after the start of the Supercomputer Centers program. NSF paved the way to link researchers to the Centers through its sponsorship of NSFNET, which augmented ARPANET as a major network backbone and eventually replaced ARPANET when ARPANET was retired in 1990. Then-Senator Gore's 1986 legislation calling for the interconnection of the Centers using fiber optic technology ultimately led the administration to respond with the High Performance Computing and Communications (HPCC) Initiative.

Among the most critical decisions that NSF made was to support the creation of "regional" or "intermediate-level" networks that would aggregate demand from the nation's universities and feed it to the NSFNET backbone. The backbone itself was initially implemented using gateways (systems used to route traffic) developed at the University of Delaware and links operating at the ARPANET speed of 56K bps. Because of rapidly increasing demand, though, NSF in 1988 selected MERIT (at the University of Michigan) to lead a cooperative agreement with MCI and IBM to develop a 1.5M bps backbone. IBM developed new routers and MCI supplied 1.5M bps circuits, and NSFNET was reborn roughly 30 times faster than its predecessor.

The regional networks quickly became the primary means by which universities and other research institutions linked to the NSFNET backbone. NSF wisely advised these networks that their seed funding would have

limited duration and they would have to become self-sustaining. Although this took longer than originally expected, most of the regional networks (such as BARNET, SURANET, JVNCNET, CICNET, NYSERNET, and so on) now have either gone into for-profit mode or have spun off for-profit operations.

Because of continued increases in demand, NSF recently re-visited the cooperative agreement with MCI, IBM and MERIT. A non-profit organization, Advanced Networks and Services (ANS), was born, and has satisfied the current demand for Internet capacity using 45M bps circuits. The name "Internet" refers to the global seamless interconnection of networks made possible by the protocols devised in the 1970s through ARPA-sponsored research -- the Internet protocols, still in use today.

A Commercial Market Emerges

By the mid-1980s there was sufficient interest in the use of Internet in the research, educational, and defense communities that it was possible to establish businesses making equipment for Internet implementation. Companies such as Cisco Systems, Proteon, and later Wellfleet (now Bay Networks) and 3Com became interested in manufacturing and selling "routers," the commercial equivalents of the "gateways" that had been built by BBN in the early ARPANET experiments. Cisco alone is already a \$1 billion business, and others seem headed rapidly toward that level.

The previous subsection noted the "privatization" of the NSF regional networks. NYSERNET (the New York State regional network) was the first to spin out a for-profit company, Performance Systems International, which now is one of the more successful Internet service providers. Other Internet providers actually began as independent entities; one of these is UUNET, which started as a private non-profit but turned for-profit and began offering an Internet service it calls ALTERNET; another is CERFNet, a for-profit operation initiated by General Atomic in 1989; a third is NEARNet, started in the Boston area and recently absorbed into a cluster of for-profit services operated as BBN Planet (recall that BBN was the original developer of the ARPANET IMP; BBN also created the Telenet service, which it sold to GTE and which subsequently became Sprintnet).

In 1988, in a conscious effort to test Federal policy on commercial use of Internet, the Corporation for National Research Initiatives approached the Federal Networking Council (actually its predecessor, the Federal Research Internet Coordinating Committee) for permission to experiment with the interconnection of MCI Mail with the Internet. An experimental electronic mail relay was built and put into operation in 1989, and shortly thereafter Compuserve, ATTMail and Sprinmail (Telemail) followed suit. Once again, a far-sighted experimental effort coupled with a wise policy choice stimulated investment by industry and expansion of the nation's infrastructure. In the past year, commercial use of the Internet has exploded.

The Roaring '90s: Privatization, and the World Wide Web

The Internet is experiencing exponential growth in the number of networks, number of hosts, and volume of traffic. NSFNET backbone traffic more than doubled annually from a terabyte per month in March 1991 to eighteen terabytes a month in November 1994. (A terabyte is a thousand billion bytes!) The number of host computers increased from 200 to 5,000,000 in the 12 years between 1983 and 1995 -- a factor of 25,000!

As 1995 unfolds, many Internet service providers have gone public and others have merged or grown by acquisition. Market valuations of these companies are impressive. America Online purchased Advanced Networks and Services for \$35 million. Microsoft supplied more than \$20 million in capital to UUNET for expansion. UUNET and PSI have gone public. MCI has unveiled a major international Internet service, as well as an information and electronic commerce service called marketplaceMCI. AT&T is expected to announce a major new service later in the year. Other major carriers such as British Telecom, France Telecom, Deutsche Telekom, Swedish Telecom, Norwegian Telecom, and Finnish Telecom, among many others, have announced Internet services. An estimated 300 service providers are in operation, ranging from very small resellers to large telecom carriers.

In an extraordinary development, the NSFNET backbone was retired at the end of April 1995, with almost no visible effects from the point of view of users (it was hard work for the Internet service providers!). A fully commercial system of backbones has been erected where a government sponsored system once existed. Indeed, the key networks that made the Internet possible (ARPANET, SATNET, PRNET and NSFNET) are now gone -- but the Internet thrives!

One of the major forces behind the exponential growth of the Internet is a variety of new capabilities in the network -- particularly directory, indexing, and searching services that help users discover information in the vast sea of the Internet. Many of these services have started as university research efforts and evolved into businesses. Examples include the Wide Area Information Service, Archie (which spawned a company called Bunyip in Canada), LYCOS from Carnegie Mellon, YAHOO from Stanford, and INFOSEEK. Aiding and stimulating these services is the recent arrival of a "killer app" for the Internet: the World Wide Web.

Developed at the European Center for Particle Research (CERN), the World Wide Web was first used in experimental form in 1989. Around 1992 it came to the attention of a young programming team at the National Center for Supercomputing Applications (NCSA) at the University of Illinois. This team developed a graphical browser for the Web, called Mosaic. In accordance with NCSA policies, this software was made widely available on the Internet for free. It took the world by storm. The excitement of being able to provide images, sound, video clips and multifont text in a hypertext system was irresistible. Between 1992 and 1995 a number of commercial versions of Web browsers and servers emerged, among them Netscape Communications, which was founded by the former chairman of Silicon Graphics, Inc., who instantly realized that the Web would dramatically magnify the utility of the Internet by replacing its rather arcane interface with something anyone could do ("point and click").

Today (May 1995) there are over 30,000 Web sites on the Internet and the number is doubling every two months. Companies that were formerly unsure about the utility of the Internet have rushed to use the Web as a means of presenting products and services. The rest of the 1990s belongs to the content providers, who will use the rapidly evolving infrastructure to bring increasingly sophisticated material to consumers.

The Future

It's risky to predict the future of something as dynamic as the Internet. It seems safe to say that we will see a continuing explosion of new services. Today, at least a dozen companies are engaged in providing electronic funds transfers on the Internet in support of electronic commerce. Other companies are exploring the provision of packetized video, videoconferencing, packetized voice (packet telephone!), and increasingly sophisticated tools for securing Internet operation for intra- and inter-corporate use.

Projections of Internet-related business range to \$50 billion at the end of the decade. While this is still small compared to the total telecommunications business (estimated at about \$300 billion today), its rapid growth and the rich evolution of new products and services suggest that the modest research investments of the federal government have paid off in myriad ways, not all of them merely monetary. There is every reason to believe that the Internet will transform education, business, government, and personal activities in ways we cannot fully fathom. Virtually none of this would have happened as rapidly, or in the same open and inclusive fashion, had not the federal government consciously provided sustained research funding and encouragement of open involvement and open standards, and then wisely stepped out of the picture as the resulting systems became self-sustaining.

The Internet is truly a global infrastructure for the 21st century -- the first really new infrastructure to develop in nearly a century.

The Author

Dr. Vinton G. Cerf is Senior Vice President of the Data Services Division of MCI Telecommunications Corporation, where he is responsible for the design of MCI's data services. He has also served as President of the Internet Society since 1992.

Cerf's career in computer communications began nearly 30 years ago, and has included stops as an independent consultant, Principal Programmer on the ARPANET Project at UCLA, faculty member in Electrical Engineering and Computer Science at Stanford University, Principal Scientist (Internet Research) at the Department of Defense Advanced Research Projects Agency, Vice President of Engineering for the MCI Digital Information Services Company, and Vice President of the Corporation for National Research Initiatives, prior to his current position. Cerf, together with Robert E. Kahn, is the co-inventor of the TCP/IP protocols, and led the Internet development effort at Stanford and ARPA from 1973 to 1982.

Cerf is a Fellow of the Association for Computing Machinery, a Fellow of the Institute of Electrical and Electronics Engineers, a Fellow of the American Association for the Advancement of Science, and a recently-elected member of the National Academy of Engineering. He is a member of the American Association for the Arts and Sciences and a 1995 recipient of the Kilby Award. He has served as chairman of the Internet Architecture Board, the ACM Special Interest Group on Computer Communications (SIGCOMM), and as chair or member of a number of National Research Council panels. He chairs the NATO Networking subcommittee of its Science Committee. Cerf is married, has two sons, and has an abiding interest in fine foods, fine wine, and mind-rotting science fiction.

[Back to Computing Research: Driving Information Technology and the Information Industry Forward](#)

Copyright 1995 by Vinton G. Cerf and the Computing Research Association. Contributions to this document by Robert E. Kahn are appreciated.

<http://cra.org/research.impact>
(better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
lazowska@cs.washington.edu

Information technology is central to our society -- economically and socially

During the past fifty years, the firms that provide our computers, communications technology, and information services, have become a vitally important underpinning to our economy and to our society -- both in terms of their own economic strength, and in terms of their productivity impacts on other industry sectors, from automobile and aircraft manufacturing, to pharmaceutical research, to overnight package delivery services.

To an increasing degree, information technology is becoming so embedded in everyday applications that it is becoming nearly invisible, hence easy to take for granted. Cellular telephones, which have become so common over the past few years, depend on highly sophisticated computing and communication technology. Advanced processors and algorithms are integral to medical diagnostic devices such as CAT scanners. Embedded microprocessors are essential components in compact disc players, video cameras, automobiles, and microwave ovens.

The development of the National Information Infrastructure, so much in the news lately, has really just begun. With proper investment, the NII holds the promise of greatly amplifying the already enormous impacts of information technology. It will extend to rural America the benefits that urban dwellers take for granted in areas such as health care, libraries, government information, cultural resources, and entertainment. It will enhance the way scientists and engineers perform the research that is so important to our nation as a whole. It will revolutionize manufacturing and commerce, and transform education.

Retaining America's leadership in information technology is vital to the nation -- to our security, to our economic competitiveness, to our employment, and to the health and well-being of our citizenry.

Back to [Computing Research: Driving Information Technology and the Information Industry Forward](#)

Copyright 1995 by Edward D. Lazowska and the Computing Research Association. Portions adapted with permission from "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure", copyright 1995 by the National Academy of Sciences. Courtesy of the National Academy Press, Washington DC.

<http://cra.org/research.impact>
(better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
lazowska@cs.washington.edu

America's leadership in information technology is the result of an extraordinarily complex and fruitful long-term partnership among government, industry, and academia

The pace of progress in information technology is so exhilarating that it's tempting to take it for granted -- to assume that it's the result of "natural forces" rather than of careful investments and a highly effective partnership in which fundamental research in computing and communications -- particularly that sponsored by the Department of Defense Advanced Research Projects Agency (through its Computing Systems Technology Office, Electronic Systems Technology Office, and Software and Intelligent Systems Technology Office) and by the National Science Foundation (through its DIRECTORATE FOR COMPUTER AND INFORMATION SCIENCE AND ENGINEERING) -- plays a key role.

It's a fact that for several decades now, the amount of computation, storage, and communication that a dollar will buy has doubled every eighteen to twenty-four months. This successive doubling -- this *exponential growth* -- is responsible for the fact that technologies such as the Internet move from near invisibility to near ubiquity in a very small number of years. Of course, the truth is that the Internet (begun as ARPAnet and evolved as NSFnet) celebrated its 25th anniversary last year -- it's been busily doubling away below most people's radar screens for most of that time. And this doubling, as will be discussed in more detail shortly, has been driven by federal investments in university research, and by the transfer, to existing and new companies, of the people and ideas that are the products of this research.

"Progress in information technology just happens," then, is one myth. Another myth arises from the enormous visibility and success of so many companies in the information technology sector: that this innovation has been driven largely by industry, and that industry can be exclusively relied upon to provide it in the future.

If you were to watch the television advertisements, you'd likely conclude that the technology underlying the nation's information infrastructure sprung forth from the minds of companies such as Microsoft and GTE. In truth, while these companies and others will play critical roles in evolving this technology and bringing it to consumers, the foundations of the technology clearly lie in federally-funded research programs which have been transferring ideas and people to the private sector for decades. As a recent study of the High Performance Computing and Communications Initiative by the National Research Council states:

"Indeed, for nearly 50 years, federal investment has helped to train the people and stimulate the ideas that have made today's computers and many of their applications possible. Federal support early in the life cycle of many ideas has advanced them from novelties, to convincing demonstrations that attract private investment, to products and services that ultimately add to the quality of U.S. life." (Page 16)

To cite a few examples of this, federal research investments played critical roles in technologies such as timesharing, computer networking, workstations, computer graphics, "windows and mouse" user interfaces, database systems, Very Large Scale Integrated (VLSI) circuit design, Reduced Instruction Set Computer (RISC) architectures, I/O subsystems based upon Redundant Arrays of Inexpensive Disks (RAID), and parallel computing.

To see how these advances fit together to create whole new possibilities for people, consider as an example the research base on which Congress's new Thomas information system rests. The revolutionary packet-switched data communication technology that forms the Internet stemmed from ARPA-funded research that started in the 1960s at sites such as UCLA, Stanford, MIT, and BBN. In the early 1980s, ARPAnet transitioned to NSFnet as the National Science Foundation took responsibility for broadening access (which involved confronting research issues of scale, in addition to simply expanding deployment). The global Internet as we know it today evolved over time as the Internet Protocol suite gained widespread acceptance due to its technological superiority and to the fact that it was in the public domain. NSF's current privatization of the Internet is an important milestone, but we must not lose sight of the fact that privatization actually has progressed steadily as the technology evolved under NSF sponsorship: for many years the backbone traffic has been increasingly carried by a commercial telecommunications company (MCI), and dozens of highly successful companies (3Com, Cisco, and Fore are examples) have been spawned along the way. Under their HPCC program responsibilities, ARPA and NSF continue to fund research that pushes the state of the art in very high speed

data communication, and to support the gigabit testbeds that explore research applications of these advanced networks. It is precisely this sort of ground-breaking research that will enable America to lead the world in the technologies that will make possible the global-scale information infrastructure of the 21st century.

The *Thomas* system also relies on many other technologies. The data query system for searching bills comes from information search and retrieval research conducted at the University of Massachusetts under ARPA sponsorship. The "Web browsers" used to retrieve information from *Thomas* -- such as the one you are using at this moment -- are based on the *Mosaic* system developed at the National Center for Supercomputing Applications, funded by NSF as part of the HPC program. (More than 1 million public domain copies of NCSA Mosaic have been obtained from NCSA, and more than 10 million commercial copies of Enhanced NCSA Mosaic have been licensed through Spyglass. NCSA started an entire new segment of the information industry with Mosaic, and populated all of the competing companies: Spyglass, Netscape, Spry, etc.) And these are only the most recent antecedents -- the roots reach back to long-term information technology research performed in a wide variety of university, industrial, and government laboratories over the past two decades.

The preceding example makes clear that the partnership responsible for America's leadership in information technology is the result of long-term, steady investment. To disrupt it now would be to damage the future in two ways: specific technology that would not be developed (in the U.S., anyway), and the destruction of a partnership that could not be reassembled "on demand" when we realized in a few years that we were falling behind.

Back to [Computing Research: Driving Information Technology and the Information Industry Forward](#)

Copyright 1995 by Edward D. Lazowska and the Computing Research Association. Portions adapted with permission from "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure", copyright 1995 by the National Academy of Sciences. Courtesy of the National Academy Press, Washington DC.

<http://cra.org/research.impact>
(better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
lazowska@cs.washington.edu

The government component of this partnership -- the highly effective public research program sponsored by ARPA and NSF -- cannot be pigeon-holed as "basic" or "applied" -- it is best characterized as "fundamental research in support of strategic directions."

As noted earlier, research in information technology is caught in the middle in a heated debate over "basic" versus "applied" research. In fact, most of the research in information technology supported by ARPA and NSF -- research supported under the High Performance Computing and Communications Initiative and under its predecessor initiatives -- occupies a complex middle ground, which can be characterized as "fundamental research in support of strategic directions."

This research is high-risk, and when it pays off it does so over an extended period -- typically 10-15 years. But it is motivated by problems that need to be solved, and when successful, it ultimately provides the people and ideas that renew existing companies and create new ones.

Many key points concerning research in information technology are carefully made in the National Research Council HPCC study. We include one table and one figure from that study here.

Table 1.1 shows that many transforming technologies that we take for granted today have their roots in the federal research program, and that while each of these research thrusts was pursued with a strategic goal in mind, the unanticipated results often were at least as significant as the anticipated ones -- a hallmark of fundamental research.

Figure 1.2, which focuses on the same technologies as Table 1.1, shows that the simple "linear model" of technology transfer -- from basic (university) to applied (industry) -- is not what happens in practice; this figure nicely illustrates the rich flow of people and ideas back and forth across the academic/industrial boundary. The figure also demonstrates that the federal research program in information technology pays off over an extended period of time: typically 10-15 years elapses from the initiation of the research to the establishment of a billion dollar industry.

While the motivation to keep the government out of very near term research is clear, an overly simplistic model of basic versus applied research would be exceedingly damaging. The ARPA and NSF research programs in information technology can be characterized as "fundamental research in support of strategic directions." It's a model that has been hugely successful.

[Back to Computing Research: Driving Information Technology and the Information Industry Forward](#)

Copyright 1995 by Edward D. Lazowska and the Computing Research Association. Portions adapted with permission from "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure", copyright 1995 by the National Academy of Sciences. Courtesy of the National Academy Press, Washington DC.

<http://cra.org/research/impact>
 (better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
 lazowska@cs.washington.edu

TABLE 1.1 Some Successes of Government Funding of Computing and Communications Research

Topic	Goal	Unanticipated Results	Today
Timesharing	Let many people use a computer, each as if it were his or her own, sharing the cost.	Because many people kept their work in one computer, they could easily share information. Reduced cost increased the diversity of users and applications.	Even personal computers are timeshared among multiple applications. Information sharing is ubiquitous; shared information lives on "servers."
Computer networking	Load-sharing among a modest number of major computers	Electronic mail; widespread sharing of software and data; local area networks (the original networks were wide-area); the interconnection of literally millions of computers	Networking has enabled worldwide communication and sharing, access to expertise wherever it exists, and commerce at our fingertips.
Workstations	Enough computing to make interactive graphic useful	Displaced most other forms of computing and terminals; led directly to personal computers and multimedia	Millions in use for science, engineering, and finance
Computer graphics	Make pictures on a computer.	"What you see is what you get" and hypermedia documents	Almost any image is possible. Realistic moving images made on computers are routinely seen on television and were used effectively in the design of the Boeing 777.
"Windows and mouse" user interface technology	Easy access to many applications and documents at once	Dramatic improvements in overall ease of use; the integration of applications (e.g., spreadsheets, word processors, and presentation graphics)	The standard way to use all computers
Very large integrated circuit design	New design methods to keep pace with integrated circuit technology	Easy access to "silicon foundries"; a renaissance in computer design	Many more schools training VLSI designers; many more companies using this technology
Reduced Instruction Set Computers (RISC)	Computers 2 to 3 times faster	Dramatic progress in the "co-design" of hardware and software, leading to significantly greater performance	Millions in use; penetration continues to increase
Redundant Arrays of Inexpensive Disks (RAID)	Faster, more reliable disk systems	RAID is more economical as well: massive data repositories ride the price/performance wave of personal computers and workstations.	Entering the mainstream for large-scale data storage; will see widespread commercial use in digital video servers

continues

TABLE 1.1—*continued*

Topic	Goal	Unanticipated Results	Today
Parallel computing	Significantly faster computing to address complex problems	Parallel desktop server system; unanticipated applications such as transaction processing, financial modeling, database mining, and knowledge discovery in data	Many computer manufacturers include parallel computing as a standard offering.
Digital libraries	Universal, multimedia (text, image, audio, video) access to all the information in large libraries; an essential need is tools for discovering and locating information	Pending development	Beginning development

A few examples

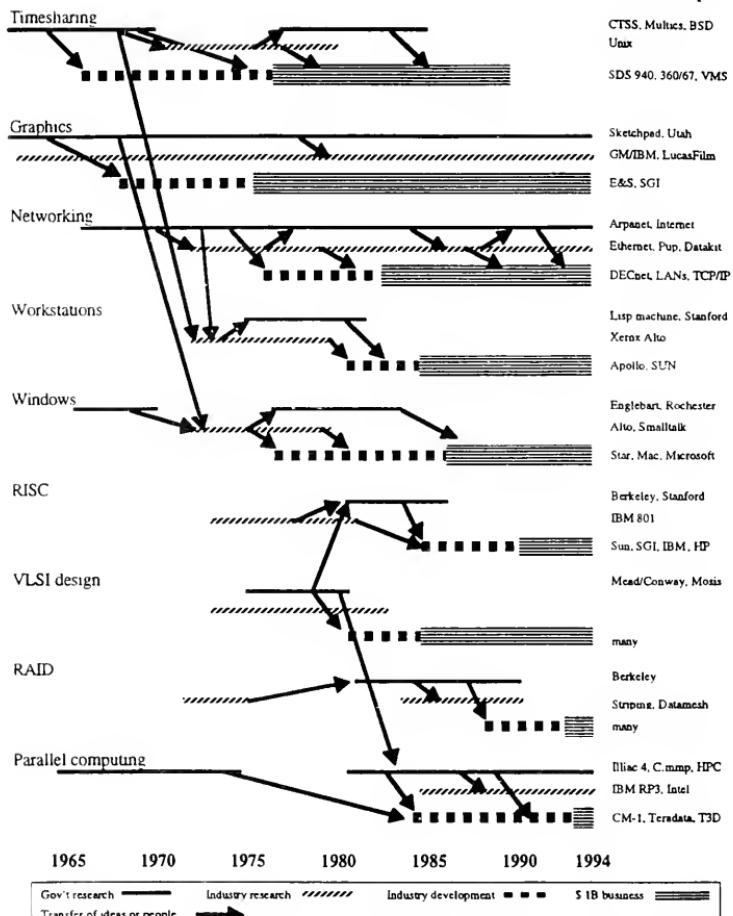


FIGURE 1.2 Government-sponsored computing research and development stimulates creation of innovative ideas and industries. Dates apply to horizontal bars, but not to arrows showing transfer of ideas and people. Table 1.1 is a companion to this figure.

Enormous possibilities still lie ahead -- in fact, the *real* information revolution is yet to come.

Predicting is difficult -- especially the future. (In information technology, under-estimation is in fact common, from Thomas J. Watson, Chairman of IBM, more than fifty years ago: "I think there is a world market for maybe five computers;" to Kenneth J. Olson, Chairman of Digital Equipment Corporation, in the late 1970s: "There is no reason anyone would want a computer in their home.")

With this caveat, there seems to be a strong bipartisan consensus at the state level, in Congress, in the Administration, in the public, and in our technical community, that the *real* information revolution -- the convergence of computing, information, and communication -- is just beginning to burst on the scene. This belief rests on two clear trends.

In the first place, as mentioned earlier, the capabilities of information technology continue to double every 18 to 24 months. The results are twofold: continuous improvements in the existing ways we use technology, and periodic breakthroughs that open up entirely new possibilities, where capabilities previously available only in the most advanced scientific and engineering laboratories become available to all. For example, in the early 1980s the cost/performance ratio of integrated circuits dropped to the point where desktop personal computers suddenly became conceivable. The way we think of and use computers has been totally transformed by that breakthrough.

Let us digress for a moment and illustrate this phenomenon of exponential growth using an example from Nicholas Negroponte's new book, *Being Digital*. It's the familiar story of the person who agrees to work for a month starting at only a penny a day, so long as her salary doubles every day. Notice two things. First, if the month has 31 working days (we'd all work weekends for this rate of pay!), the total salary collected is more than \$20 million. Second, half the salary is collected on the last day. Similarly for information technology, assuming that we invest to sustain past rates of progress: the absolute performance is staggering, and tomorrow's payoff will be even bigger than today's. It is critical that the United States reap these benefits.

In the second place, the integration of computers with communications and other information technologies to form whole new forms of distributed information systems and services has just begun. The NII will therefore amplify the already phenomenal increase in performance. That is, we not only will have more powerful personal computers, but these computers will be connected to a world-wide network of services and resources.

The real revolution is this digital convergence -- the computer as an information access device. This revolution, as noted earlier, will extend to rural America the benefits that urban dwellers take for granted in areas such as health care, libraries, government information, cultural resources, and entertainment. It will enhance the way scientists and engineers perform the research that is so important to our nation as a whole. It will revolutionize manufacturing and commerce, and transform education.

America can, and America must, lead this revolution.

[Back to Computing Research: Driving Information Technology and the Information Industry Forward](#)

Copyright 1995 by [Edward D. Lazowska](#) and the [Computing Research Association](#). Portions adapted with permission from "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure", copyright 1995 by the [National Academy of Sciences](#). Courtesy of the [National Academy Press](#), Washington DC.

<http://cra.org/research/impact>
 better performance from <http://www.cs.washington.edu/homes/lazowska/cra/>
 lazowska@cs.washington.edu

Continued federal investment is *critical* to realizing this promise -- to continuing our nation's momentum and to retaining our lead.

To realize the promise of information technology -- to continue to drive this exponential progress -- requires continued investment in long-term research in computing and communications.

It is nearly impossible to predict where and when the next major breakthrough will occur. However, one can examine objectives and derive ideas of where research investments could be strategically made. One year ago the Computing Research Association, in cooperation with a variety of government, industry, and academic individuals and organizations, sponsored a meeting to outline a research agenda aimed at realizing the vision of the National Information Infrastructure. The results were published in the report *R&D for the NII: Technical Challenges*. More recently, the National Science and Technology Council's Committee on Information and Communications, chaired by Dr. Anita K. Jones, Defense Director of Research and Engineering, produced a Strategic Implementation Plan through a process co-chaired by Dr. Paul Young, NSF's Assistant Director for Computer and Information Science and Engineering, and Mr. John Toole, then head of ARPA's Computer Systems Technology Office and now Director of the National Coordination Office for HPCC. The plan, *America in the Age of Information*, identified six Strategic Focus Areas "to focus fundamental information and communications research and to accelerate development in ways that are responsive to NSTC's overarching goals, agency mission goals, and our Nation's long term economic and defense needs." These Strategic Focus Areas are global-scale information infrastructure technologies, high performance / scalable systems, high confidence systems, virtual environments, user-centered interfaces and tools, and human resources and education.

A natural question, in light of the history of success and the promise for the future, is "Why shouldn't industry take responsibility for funding information technology research?" There are two answers to this question, both developed in detail in the National Research Council HPCCI study, from which we will briefly quote:

"First, ... few companies will invest for a payoff that is 10 years away, and even a company that does make a discovery may postpone using it. The vitality of the information technology industry depends heavily on new companies, but new companies cannot easily afford to do research; furthermore, industry in general is doing less research now than in the recent past. But because today's sales are based on yesterday's research, investment in innovation must go forward so that the nation's information industry can continue to thrive.

"Second, it is hard to predict which new ideas and approaches will succeed. The exact course of exploratory research cannot be planned in advance, and its progress cannot be measured precisely in the short term. The purpose of publicly funded research is to advance knowledge and create new opportunities that industry can exploit in the medium and long term, not to determine how the market should develop." (Page 4)

"The government-supported research program (on the order of \$1 billion for information technology R&D) is small compared to industrial R&D (on the order of \$20 billion), but it constitutes a significant portion of the research component, and it is a critical factor because it supports the exploratory work that is difficult for industry to afford, allows the pursuit of ideas that may lead to success in unexpected ways, and nourishes the industry of the future, creating jobs and benefits for ourselves and our children. The industrial R&D investment, though larger in dollars, is different in nature: it focuses on the near term -- increasingly so, as noted earlier -- and is thus vulnerable to major opportunity costs." (Page 24)

The NRC HPCC study also includes the following remarkable quotation from Alexander Hamilton, from the 1791 *Report on Manufactures*:

"Industry, if left to itself, will naturally find its way to the most useful and profitable employment. Whence it is inferred that manufacturers, without the aid of government, will grow up as soon and as fast as the natural state of things and the interest of the community may require.

"Against the solidity of this hypothesis ... very cogent reasons may be offered ... [including] the strong influence of habit; the spirit of imitation; the fear of want of success in untried enterprises; [and] the intrinsic difficulties incident to first essays towards [competition with established foreign players]: the bounties, premiums, and other artificial encouragements with which foreign nations second the exertions of their own citizens ..."

"To produce the desirable changes as early as may be expedient may therefore require the incitement and patronage of government."

The information technology research supported by the Department of Defense Advanced Research Projects Agency and by the National Science Foundation exactly fulfills these objectives. In the 1960s ARPA and NSF started programs that were instrumental in developing the disciplines of computer science and computer engineering in universities across the nation. ARPA's efforts tended to be focused on large investments at a modest number of sites. NSF's efforts tended to involve smaller investments in a larger number of areas. These two models complemented one another perfectly, giving the nation a balanced portfolio of research in information technology which led directly to America's undisputed world leadership in this field. In turn, the contribution of this field to the nation's economic, social, and military security has been quite literally incalculable.

We cannot rest, though. To quote one final time from the NRC HPCC study:

"Our lead in information technology is fragile, and it will slip away if we fail to adapt. Leadership has often shifted in a few product generations, and a generation in the information industry can be less than 2 years. As a nation we must continue to foster the flow of fresh ideas and trained minds that have enabled the U.S. information technology enterprise as a whole to remain strong despite the fate of individual companies." (Page 16)

Back to Computing Research: Driving Information Technology and the Information Industry Forward

Copyright 1995 by Edward D. Lazowska and the Computing Research Association. Portions adapted with permission from "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure", copyright 1995 by the National Academy of Sciences. Courtesy of the National Academy Press, Washington DC.

<http://cra.org/research/impact>
 (better performance from <http://www.cs.washington.edu/homes/lazowska/cra>)
 lazowska@cs.washington.edu

Related material

Testimony

- Oral testimony of the Computing Research Association before the House Committee on Appropriations concerning the FY 1996 NSF appropriation, April 5, 1995
- Congressional staff briefing by the Computer Systems Policy Project concerning federal support of university research, April 19, 1995
- Written testimony of the National Coordination Office for HPCC before the Senate Committee on Science, Technology, and Space concerning the HPCC program, May 4, 1995

Letters and Articles

- A Moment of Truth for America (An open letter to Congress from the executives of some of America's leading technology companies, May 1995)
- A letter to Senator Robert Dole and Congressman Robert Walker from the Computer Systems Policy Project, May 24, 1995
- A letter to the Congressional leadership from Microsoft Corporation, June 8, 1995
- Research Scope Narrows: R&D Cuts, Corporate Focus on Near-Term Payoffs Could Test Health of U.S. Economy (Knight-Ridder newspaper article, April 23, 1995)

Reports

- High Performance Computing and Communications: Foundation for America's Information Future (FY 1996 HPCC "Blue Book")
- Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure (National Research Council HPCCI study)
- America in the Age of Information (Strategic Implementation Plan of the Committee on Information and Communications of the National Science and Technology Council)
- R&D for the NII: Technical Challenges (Report of a workshop)

Organizations

- National Coordination Office for HPCC
- Department of Defense Advanced Research Projects Agency
 - Computing Systems Technology Office
 - Electronic Systems Technology Office
 - Software and Intelligent Systems Technology Office
- National Science Foundation
 - Directorate for Computer and Information Science and Engineering
- Computing Research Association

Back to [Computing Research: Driving Information Technology and the Information Industry Forward](#)

A MOMENT OF TRUTH FOR AMERICA.

Imagine life without polio vaccines and heart pacemakers. Or digital computers. Or municipal water purification systems. Or space-based weather forecasting. Or advanced cancer therapies. Or jet airliners. Or disease-resistant grains and vegetables. Or cardiopulmonary resuscitation (CPR).

We take for granted these and thousands of other technological breakthroughs that have made American society the most advanced in history. They have made our economy more competitive, created millions of jobs, and underpinned our entire standard of living. They have vastly improved our health and extended our life span. In a very real sense, they epitomize the American Dream.

But these breakthroughs didn't just happen. They are the products of a long-standing partnership that has, as a matter of national policy, fostered the discovery and development of new technologies. For many years, Administrations of both parties, working with Congress, have consistently supported university research programs as a vital investment in our country's future. Industry has played an equally critical role, carefully shepherding these new technologies into the marketplace.

This partnership — the research and educational assets of American universities, the financial support of the federal government and the real-world products developed by industry — has been a critical factor in maintaining the nation's technological leadership through much of the 20th century.

Just as important, university research has also

helped prepare and train the engineers, scientists and technicians in industry whose discipline and skill have made technological breakthroughs possible. It has sparked innovation and prudent risk-taking. And as a result of the opportunity afforded such skilled workers in our technologically advanced economy, many disadvantaged young people have used high-tech jobs as a "stepping stone" to more productive and satisfying lives.

Unfortunately, today America's technological prowess is severely threatened. As the federal government undergoes downsizing, there is pressure for crucial university research to be slashed.

University research makes a tempting target because many people aren't aware of the critical role it plays. It can take years of intense research before technologies emerge that can "make it" in the marketplace. History has shown that it is federally sponsored research that provides the truly "patient" capital needed to carry out basic research and create an environment for the inspired risk-taking that is essential to technological discovery. Often these advances have no immediate practical usability but open "technology windows" that can be pursued until viable applications emerge. Such was the case with pioneering university research done on earthquakes in the 1920s, which led over time to the modern science of seismology and the design of structures that better withstand earthquake forces.

Today, we, the undersigned — executives of some of America's leading technology companies —

believe that our country's future economic and social well-being stands astride a similarly ominous "fault line." We can personally attest that large and small companies in America, established and entrepreneurial, all depend on two products of our research universities: new technologies and well-educated scientists and engineers.

Technological leadership, by its very nature, is ephemeral. At one point in their histories, all the great civilizations — Egypt, China, Greece, Rome — held the temporal "state of the art" in their hands. Each allowed their advantage to wither away, and as the civilization slipped from technological leadership, it also surrendered international political leadership.

For all these reasons, it is essential that the federal government continue its traditional role as funder of both basic and applied research in the university environment. If we want to keep the American Dream intact, we need to preserve the partnership that has long sustained it. As we reach the final years of the century, we must acknowledge that we face a moment of truth:

Will we nurture that very special innovative environment that has made this "the American century"? Or will we follow the other great civilizations and yield our leadership to bolder, more confident nations? As the Congress makes its decisions on university research, let there be no mistake: We are determining the 21st century today.

W.W. Allen
W. Wayne Allen
Chairman & CEO
Phillips Petroleum Company

Norman R. Augustine
Norman R. Augustine
President
Lockheed Martin Corporation

Robert L. Culp
Robert L. Culp
Chairman & CEO
Chrysler Corporation

George Fisher
George H. C. Fisher
Chairman, President & CEO
Eastman Kodak Company

Robert W. Galvin
Robert W. Galvin
Chairman, Executive Committee
Motorola, Incorporated

Joseph T. Gorman
Joseph T. Gorman
Chairman & CEO
TKW, Incorporated

Gerald Greenwald
Gerald Greenwald
Chairman & CEO
United Airlines

George H. Heilmann
George H. Heilmann
President & CEO
Bellcore

John McDonnell
John McDonnell
Chairman
McDonnell Douglas Corporation

Randall L. Tobias
Randall L. Tobias
Chairman & CEO
Eli Lilly and Company

P. Roy Vagelos, M.D.
P. Roy Vagelos, M.D.
Former Chairman & CEO
Merck & Company, Incorporated

John F. Welch
John F. Welch
Chairman & CEO
General Electric Company

Edgar S. Woolard, Jr.
Edgar S. Woolard, Jr.
Chairman & CEO
E.I. DuPont deNemours and
Company



Michael Spindler Apple

Robert E. Allen AT&T

Eckhard Pfeiffer Compaq

John F. Carlson Cray Research

Ronald L. Skates Data General

Robert B. Palmer Digital Equipment

Lewis E. Platt Hewlett-Packard

Louis V. Gerstner, Jr. IBM

Edward R. McCracken Silicon Graphics

William E. Foster Stratus Computer, Inc.

Scott G. McNealy Sun Microsystems

James G. Treybig Tandem

James A. Unruh Unisys

May 24, 1995

The Honorable Bob Dole
U.S. Senate
Washington, D.C. 20510

COPY

Dear Senator Dole:

On behalf of the Computer Systems Policy Project (CSPP), an affiliation of Chief Executive Officers of thirteen American computer systems companies, we urge you to maintain federal funding for university research in information technology. The United States has enjoyed unparalleled success in the invention and application of information technology. This story is effectively chronicled in a recent National Research Council report, "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure." The foundation for this incredible innovation engine has been the pioneering work done in our research universities. It would be a tragedy, in this era of increased global competition, if the United States cut back its support and funding of such strategic research.

We are not asking for funding for our own companies, however, we believe that the U.S. would greatly benefit from federal dollars spent on information technology research performed at universities. In fact, federal support frequently triggers the creation of companies and businesses that often become our competitors. However, we all benefit significantly as innovation and competition within this country sets the world wide pace for progress in information technology.

While we appreciate the need to take action to reduce the national budget deficit, we strongly believe that reducing funding for information technology research at U.S. universities would be counterproductive. The U.S. computer systems industry relies on the university research supported by the National Science Foundation, the Department of Defense, and other agencies to construct the foundation from which to build the technology advancements that have made the U.S. a global leader in both market share and technological innovation.

In particular, reducing funds for the High Performance Computing and Communications Initiative and other National Information Infrastructure (NII) related research would seriously jeopardize the progress that this nation's researchers have been making in the critical areas of computing, software, and communications. The research supported by these federally funded programs is especially

KENNETH R. KAY, Executive Director
1001 G Street, NW, Suite 900 East
Washington, DC 20001
202 393-0220 (Fax) 202 393-5510

May 24, 1995
Page 2

important in helping to accelerate the deployment by U.S. industry of an enhanced NII. An advanced NII will link institutions, individuals, and information resources nationwide to enable new applications in health care, education, manufacturing, and access to government information and services. University researchers are working to advance many of the technologies that are critical underpinnings of the NII, such as storage technologies, technologies to protect information security, ensure privacy and confidentiality, protect intellectual property, and authenticate information sources.

As you prepare to consider the 1996 budget for the coming year, we strongly urge you to sustain funding for university information technology research at its current level.

Sincerely,

David Nagel
Apple Computer, Inc.

Joel Birnbaum
Hewlett-Packard Company

Tom A. Mays
AT&T Corporation

James C. McGroddy
IBM Corporation

Robert Stearns
Compaq Computer Corporation

Forest Baskett
Silicon Graphics

Steve Nelson
Cray Research, Inc.

Stephen C. Kiely
Stratus Computer, Inc.

J. Thomas West
Data General Corporation

Ivan Sutherland
Sun Microsystems, Inc.

Samuel H. Fuller
Digital Equipment Corporation

Kurt L. Friedrich
Tandem Computers, Inc.
Ron Bell
Unisys Corporation

Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399

Tel 206 882 8080
Telex 160530
Fax 206 936 7329

936 1222

Microsoft

June 8, 1995

The Honorable Robert Dole
United States Senate
Washington, D.C. 20510

Dear Senator Dole:

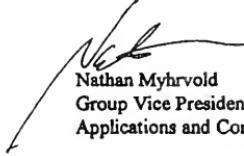
I am writing to express my strong support for the High Performance Computing and Communications Initiative (HPCCI). HPCCI -- through the Department of Defense's Advanced Research Projects Agency (ARPA) and the National Science Foundation's Directorate for Computer and Information Science and Engineering -- is responsible for much of the nation's fundamental research in the critical field of information technology.

Information technology is vital to our society. The United States is the recognized leader in information technology globally. Information technology companies are a cornerstone of our economy and one of its most rapidly growing sectors. Much of this country's success in the field comes from the kinds of long-term partnerships among government, academia and industry that are fostered by the HPCCI.

Dismantling the High Performance Computing and Communications Initiative and the ARPA and NSF information technology research programs would have serious long-term implications for U.S. industry and U.S. competitiveness. We cannot afford to lose one of the few edges that we have.

I urge you to support continued funding of this valuable initiative.

Sincerely,



Nathan Myhrvold
Group Vice President
Applications and Content

Mr. SCHIFF. Thank you very much.
Dr. Ostriker I believe you are next.

STATEMENT OF DR. JEREMIAH P. OSTRIKER, CHARLES A. YOUNG PROFESSOR OF ASTRONOMY, PROVOST, PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY

Dr. OSTRIKER. Thank you, Mr. Chairman and members of the Committee and staff. I am glad to be here today and to describe one part of the HPCC Initiative and that is the Grand Challenge program and in particular a scientific component, the efforts in astronomy and cosmology.

Representative Walker alluded to the "fundamental scientific research that pushes the envelope." Well, that is what I hope exactly what my own field is.

There are two academics who are speaking to you today and they are in fields that look from a distance very much like one another. This is computer science and I am scientific computing. Close up they look a little bit further apart.

So, let me tell you a little bit about who I am and what I do.

I have been a professor at Princeton University of theoretical astrophysics for 30 years, and have worked in a variety of fields, pulsars, quasars, all those strange things that you have heard about. I am one of the people who discovered dark matter, working in cosmology.

Until a few years ago, I would have been happy to testify on the other side of this issue. I had the firm conviction that supercomputing was not—the scientific aspects of it—was not only useless, it was positively pernicious, that it had been counterproductive, and that fairly large amounts of resources had gone into it without much results.

The reason was that it was very, very specialized. Only a very small number of people could do it. The vast bulk of the results that came out of scientific computing—I am not talking about all the other aspects of science—came out of relatively small scale efforts, the equivalent of what today we would call workstations. The often very able people who went off to work on supercomputers, so to speak, disappeared into an intellectual black hole. They did all kinds of things which were exciting to themselves, but the rest of the world and their scientific colleagues never heard about it.

I have turned in a sense and I am now preaching to you with the enthusiasm of the convert. I am a principal investigator of one of these HPCC Grand Challenges, and I feel that this new technology is producing really revolutionary new science, a new kind of science.

So, the first question is, what is the reason for the change in my attitude? I wrote it up in the shortest of all of the handouts. It is a one-page one.

I noticed in listening to the other testimony and in reading the prepared material how parallel were the descriptions. There is an essential point here which is a technical point but it is very simple to understand. Given the computing power that was available when I was in graduate school, you could do one-dimensional problems. You find out something which was spherical or cylindrical or something very simple. Then a decade later you could do two-di-

dimensional problems. Well, then a decade later, which is now or a few years ago, you could start doing three-dimensional problems. Well, everything we see around us, all of these objects are three-dimensional objects. There are not any more dimensions. There are three dimensions and time. So, once you can do three-dimensional problems, you can work in the real world.

Well, the first things that were done in this area were done with a very, very low resolution. The first ones that were done in cosmology were 30 by 30, and you can imagine a picture of a face which was 30 by 30 pic cells, you could barely recognize that it was a face.

When I got into the field, it was about 300 by 300 which is about the resolution of U.S. television, which is quite adequate for most people but is a little bit grainy by current standards. What you can now do is about 3,000 by 3,000 which is pretty close to what the human eye can see and is the resolution of good photographs. So, that is the state of the art. By this time if you can simulate the world in a computer at the level of resolution which is what we can see and perceive, that is high enough resolution to be talking about things of real interest, whether it is oil field reservoirs or it is cosmology and pictures of the sky.

What it has done is led to the birth of a really new field of science. In the past you could break science roughly into two categories, a descriptive one which described the world as it really was—and I gave the example of descriptive astronomy or even astrology which is just plotting the motions of the planets through the constellations—or theoretical science, which I gave the examples of, Newton's laws and Einstein's law which are mathematically based and enable you to understand the origins of the world we see.

But there was always an enormous gap between these two fields because the physical laws that we understood, we could only use them to solve the simplest possible problems. For example, the motion of an artillery shell, that could be done, or the strength of a bridge. Problems like that could be solved from first principles in the 19th century.

If you ask complicated three-dimensional problems like what is the drag on a ship or what is the path of a hurricane, we absolutely knew the equations that we would have to solve to determine this. It is just that we could not solve them. It was impossible. It was beyond our capabilities.

Well, that has changed. The new machines enable us to solve those problems, and as of a few years ago, we could solve them badly. Now we can solve them pretty well, and at the current rate of development, we will be able to solve them very well. I think whole fields of endeavor will fall to this new discipline of numerical simulations.

Let me just talk about my own field of cosmology briefly and then I will try to sum up.

In that, that was very much like theology for an awfully long time, where people had wonderful theories which you could not really in any sense compare to the actual observed universe that existed around us. So, you could not say, if this theory is true, should that star be there or planets or galaxies? So, theories were

judged, so to speak, on the basis of their beauty rather than their correspondence to reality, and the simple reason was that you could not conclude what the theories really implied for the real world.

That has changed and it is simply due to the extraordinary progress in hardware and software that we have had detailed before us.

So, now one can take theories which I had not realized myself (a few years ago) but are completely predictive theories, things that you have heard defined by things like the cold dark matter theory and you and I might have thought them as essentially theological. I realized that they were absolutely predictive theories, and if they were right, they would tell you just what the real world should be like. Now we can put those theories in the computer, turn the crank, so to speak, and see what happens and see, does it predict the universe that is seen by the Hubble space telescope and does it predict the universe that is seen by the very large array?

By the way, these are very high resolution pictures of the world and you have to be able to generate in the computer equally high resolution pictures to see whether or not these two things are the same. Well, we can now do it.

Have we found out what the correct theory is? No. I cannot claim that I can say how the universe started out being almost completely uniform and fragmented into clusters of galaxies, galaxies, planets, stars, you, and me. But it is good enough so that we can now actually rule out some theories. There are some theories that people thought could have been right which we can now show would predict a universe very different from the one we live in.

Well, what has this cost us? On the one hand, when you look at the price of individual machines, it seems like a lot of money. On the other hand, if you look at the fraction of the resources—let us say at a university like Princeton, Carnegie Mellon, Cal Tech—which goes into various components of computing, the very high end is a tiny fraction. I made some recent estimates and it is about 55 percent of the dollars in a place like Princeton, now wearing my provost hat, are spent on things like PC's or Mac's, about 40 percent on workstations, and maybe 5 percent on the high end with maybe 1 percent on the very high end that we are talking about today. And I mean indirectly, even taking into account that the Government may be “giving” us these, that the actual value of the resources being used is a small fraction of what goes into computation.

But it is this very small bit that is right on the edge of technology which both feeds on the way down producing the algorithms that drive everything else, but also it has a vacuum really. It attracts many of the best people because of the excitement of the problems. A student that I had working with me in getting his Ph.D. last year was just the best graduate student in Princeton University and the same for the year before. This is extraordinary that in one small area the best graduate student in Princeton University would be in the same field 2 years in a row, and it has just attracted very, very able people from all over the world who then go out and do all kinds of things. Some of them do science. Some of them do other things.

So, in summary, I would say that the HPCC Initiative and the Grand Challenge part of it has worked. It has harnessed some of the best minds in the world to some of the most exciting and fundamental problems that we know of, and whereas it has not solved all the problems, it has certainly solved some of them and restricted the range of possible solutions so we know an awful lot more than we did and can utilize the Government's investment in many of the very large experimental programs much better.

I will be happy to answer any questions as may come up.

[The prepared statement of Dr. Ostriker follows:]

Testimony before the Subcommittee on Basic Research
Hearing on: High Performance Computing and Communications Program (HPCC)
Tuesday, October 31, 1995

Presented by: Dr. Jeremiah P. Ostriker, Provost; Professor of Astrophysical Sciences
Princeton University, Princeton, NJ

HIGH PERFORMANCE COMPUTING AND THE SCIENCE OF NUMERICAL SIMULATIONS

High performance computing in general and the United States HPCC program in particular have enabled the creation of a new branch of science. Previously, we had one branch of experimental/observational science, which aimed at a detailed and complete *description* of nature, and another branch of theoretical science, which aimed at precise understanding of the physical laws behind the observed phenomena. In my own field, descriptive astronomy or even astrology (!) are on the first branch, with Newton's and Einstein's laws on the second branch.

However, there are vast domains of phenomena in which we know the laws which are relevant but cannot calculate the consequences of those laws. For example, it has been possible for over a century to use the laws of physics to correctly predict the path of an artillery shell or the carrying capacity of a bridge. But, it has not been possible to use those same laws to compute the best shape for a ship or the path of a hurricane or, to take an example from my subject, the manner in which a star like our Sun was formed. We could not solve these problems even if we thoroughly understood all of the physical principles needed to solve them, simply because the equations became too difficult, too complicated to solve.

When I was in graduate school (~1965), one-dimensional problems could be solved on computers (*e.g.* the structure of a spherical star); then a decade later (~1975) two-dimensional problems (*e.g.* the motion of a single water wave) could be solved, and then by 1985, we were beginning to be able to solve real three-dimensional problems (*e.g.* weather). But, there are, for practical purposes, only three-dimensions in nature (plus time), so the work in the last decade has gone into solving 3-D problems *properly*.

When three-dimensional astronomical simulations began a decade ago, the resolution was $30 \times 30 \times 30$, which, if you think in terms of a 30×30 picture, is so crude that a picture of a face would be barely recognized. A few years ago the resolution was 300×300 , about the same as U.S. television, but the current best resolution is about $3,000 \times 3,000$, the same as the best photographs, providing a level of detail so that one can compare the computer simulations to the real world in a meaningful way.

Cosmology was in the past more a religion than a science, as it was impossible to really confront the various theories with evidence from the real world. Now, we can take a theory like

the Cold Dark Matter Theory for the origin for structure in the universe, and, realizing that it is, in principle, a perfectly predictive theory, put it into the computer and ask what it would predict. Then, we can take these predictions in the form of pictures (or numbers) and compare them with the pictures (or numbers) being generated by the Hubble Space Telescope, the Very Large Radio Array in New Mexico, or any of the other exciting new scientific instruments recently put into place to study the heavens. Then, we can ask the simple question, "Does the computed picture look like the real world?" And so we can prove that some cosmological theories are false. And, we have done this. We can even hope to tell in the next few years which are the best candidates for the true theory.

An important element in the transformation of this science is that now the programming languages and, in fact, the computer chips themselves, may be the same at the lowest level of a PC and the highest level of the greatest supercomputer. Thus, students can pass seamlessly up the ladder from personal machines to departmental and mid-size machines to the largest ones available, those needed to attack the major scientific problems.

The fact that we have in the United States ample opportunity for young people to work at every level of this pyramid with the hope of reaching the apex has attracted some of the brightest young people in the U.S. (and in the world) to this program. The current rapid rate of technological change in this field makes me confident in predicting that one problem after another will fall to the advance of this new discipline, if only we can stay the course. The HPCC initiative has been a great success to date, and it is certainly worth pursuing for another decade.

Mr. SCHIFF. Thank you, Dr. Ostriker.
Dr. Baskett.

STATEMENT OF DR. FOREST BASKETT, CHIEF TECHNOLOGY OFFICER AND SENIOR VICE PRESIDENT OF RESEARCH AND DEVELOPMENT, SILICON GRAPHICS, INC., REPRESENTING THE COMPUTER SYSTEM POLICY PROJECT

Dr. BASKETT. I am Forest Baskett. I am Senior Vice President of Research and Development at Silicon Graphics. I am also Chairman of the Chief Technologists Committee of the Computer Systems Policy Project, the CSPP, and I am here in that capacity today.

The CSPP is a group of chief executive officers from the 13 leading American computer systems companies. On their behalf, as well as my own, I am pleased to be here today in support of the high performance computing and communications program.

For starters, let me just first say that there are really three major points that I would really like to leave you with today if you hear nothing else, and let me tell you what those are up front.

The HPCC Program really fosters U.S. innovation. HPCC funding is driving innovation in the information technology industry, and that innovation is happening in the United States and that innovation is leading the world. Without HPCC funding, innovation will certainly continue but at a slower rate and it will be worldwide. The playing field today where the U.S. is leading would become level.

Secondly, if the HPCC Program were to be severely cut, it is highly unlikely that the U.S. information technology industry would fill the void. The competitive nature of the current information technology industry means that companies that continue to survive do not and will not fund significant amounts of long-range research. The long-range nature of HPCC funding is why we are leading the world and we of the CSPP see no substitute for it.

Thirdly, products from the information technology industry are pervasive. Key U.S. industries benefit from information technology innovations. When our products become 100 times more powerful every 7 years, we change almost every industry, manufacturing industries, service industries, financial industries, biotechnology industries, everything.

Silicon Graphics today is a multibillion dollar company that is part of the U.S. information technology industry. Some estimate that the information technology industry generates about 16 percent of our gross domestic product and a much larger percent of our total dollar volume of exports to the rest of the world. Computer systems companies like the members of CSPP are not the whole of the information technology industry, but computer systems are a key part of communications systems, a key part of information services, and certainly a key part of the software industry. As such, we are keenly interested in the High Performance Computing and Communications Program.

Silicon Graphics itself traces its existence back to federally funded science and technology. An ARPA program in very large scale integrated circuits had a major component at Stanford University. As part of that program, Professor Jim Clark realized that it would

be possible to do on a single silicon chip what had previously only been possible by a room full of traditional electronics, and that was to process and manipulate three-dimensional models in real time. The ARPA program made it possible for Jim to demonstrate that realization. That demonstration was just one of many that the ARPA program made possible, and it was just one of many that then convinced the traditional capital markets that a real technological innovation had been achieved and that a new industry was feasible based on that innovation. So, Silicon Graphics was founded by Professor Clark and several of his graduate students as a result of the ARPA-funded research they had performed at Stanford.

Many other companies were founded and several are as successful as Silicon Graphics. Dr. Sutherland's chart I think indicates several of those companies.

Since that time, we at Silicon Graphics have continued to work with Stanford and other universities, and we have continued to find new innovations that have continued to fuel our own growth and our leadership in world markets.

In fact, the members of CSPP believe that Federal support of science and technology has been the innovation engine behind our industry. The discoveries that have been made and that are continuing to be made that have kept us so far ahead of the world in this industry are most often the results of work initiated as part of Federal research programs such as the ARPA program that I mentioned. Today many of those research program in information technology are part of the HPCC Program.

On the other hand, we recognize that the days of continued total growth in that support are over. Today's budget realities mean that we must be more thoughtful and more careful in our choices of what programs deserve support. So, today I would like to describe some of the principles that we believe should help guide those choices.

Much of the innovation in our industry is made possible by the continued improvement in silicon device fabrication capabilities. As we look into the future, we are confident in predicting that the 60 percent per year improvement in silicon device densities that we have been seeing for the last 30 years will continue for at least another 10 years. The 50 percent per year improvement in silicon device speeds will continue for at least another 10 years. The power of new innovations derives from this power base which is the product of the device speeds and the device densities. So, we can easily expect to be able to do 100 times as much with a silicon device in 7 years than we are able to do today. With this kind of change in the power of our tools, our accomplishments are only limited by our imaginations and our opportunities for being imaginative.

The HPCC Program provides many of those opportunities. Without those opportunities, we spend most of our time simply trying to keep up with the technology in existing systems. Existing systems do get faster. Existing systems do get less expensive. But existing systems do not create whole new sub-industries and industries. Let me try to imagine a few examples.

Will we be able to change the growth prospects for the commercial airline industry by inventing technologies that would make vir-

tual meetings as good as or maybe even better than real meetings? It is possible. We have seen indications that it is really possible.

Could we invent a system that could schedule taxis for a city and its suburbs that could work so well and so efficiently and be so integrated with mass transit that people would prefer it to private autos for most daily business travel? Would such a system change the nature of our cities? Well, this computational problem is only about 1,000 times more difficult than the airline scheduling problem. So, it is possible. I could go on except for my own limited imagination.

The only other industry that even comes close to the information technology industry in its fundamental ability to innovate I believe is the biotechnology industry.

Now, our industry today is a thriving and healthy one. It is vigorous. It is exciting. The U.S. part of it is clearly leading the world, but it is also a very competitive industry. That competition is part of why it is so vigorous. That competition has expanded our industry by bringing more new innovations to market more quickly and more cheaply.

On the other hand, while that competition has expanded our industry, it has also reduced our internal spending on R&D and caused most of that internal spending to be on development, not research. Many years ago when there was one dominant computer company and one dominant communications company, those companies had research enterprises that were the envy of the world. Those research enterprises existed for at least two reasons. One was that partly because of the dominance of those companies, there was not a sufficient base of outside research to draw on to fuel new development. The other was that those companies' dominant position meant that they could afford to support those research enterprises without endangering the existence of the company.

Today a big internal, long-range research program by a company is a burden that can easily kill the company in the competitive marketplace before any benefits of that research can be realized. We have become increasingly dependent on both university research and university-trained professionals to fuel our continued growth and our continued leadership and our continued prosperity. The large private research enterprises of the past have been significantly scaled back in size and redirected at short-term goals.

Federal support of research is not what has caused this reduction in internal R&D and its focus on development; competition has. Federal support of research has simply allowed us to continue to lead the world in this worldwide competitive environment. If Federal support were to disappear, the playing field would become level. The rate of improvements in existing technologies and the rate of introductions of new innovations and new technologies and new industries would decrease. U.S. leadership would be in danger.

Our companies for the most part are not looking to the Federal Government to support our internal R&D programs. In fact, the Federal grants and contracts process is too slow for our fast-moving industry. The Federal reporting process for those grants and contracts is burdensome, and the traditional capital markets are fast and efficient. Debt capital, equity capital, and venture capital are all available for worthy, short-term projects. It is the long-term

projects that are in need if we are to maintain our world leadership.

A significant issue for computer system companies and, I believe, for all information technology companies is a current discussion of basic and applied research. We have difficulty with some of this discussion because of our inability to understand any reasonable distinction between the two. We are most sympathetic with Nobel laureate Lord Porter who said, "There are two kinds of research—applied research and not-yet-applied research." We do understand the distinction between research and development, however, but we worry that an artificial distinction between basic and applied research will endanger the research programs at our world-class system of schools of engineering. The research programs at these schools of engineering have been a major source of innovation for our industry and certainly a major source of fresh, new, innovative minds trained in the disciplines we most need.

Now, there are several features of the HPCC Program that we wholeheartedly support. Let me just list them briefly but encourage you to consider them more thoughtfully.

First, merit review for Federal support.

Secondly, the fact that support is for programs not for organizations or institutes or institutions.

Thirdly, support that generates both research and training.

Fourth, the multi-agency coordination and collaboration. Here the HPCC Program is superb.

Support for both science and engineering research, especially at universities.

And finally, support for both theory and experiment. The experiments test the theories and train the students, and the experiments provide the demonstrations that ensure that technology transfer happens for the real innovations.

Significantly reducing the HPCC Program and other national information infrastructure related research would seriously jeopardize the progress that this Nation's researchers have been making in critical areas of computing, software, and communications. The HPCC Program has been instrumental in expanding our knowledge about how to use high performance computers and high speed networking to solve a wide range of scientific and engineering problems.

It is also providing the research foundation for the national information infrastructure which will bring the benefits of information technologies to all Americans through applications in education, health care, access to Government information and services, and digital libraries. University researchers are working to advance many of the technologies that are the critical underpinnings of the NII, such as storage technologies, technologies to protect information security, ensure privacy and confidentiality, protect intellectual property, and authenticate information sources.

Many industrial sectors are using the technologies advanced through the HPCC Program to improve their manufacturing efficiencies, shorten production times, reduce costs, and develop products that improve the quality of life. There are many examples. The auto industry is using high performance computers to design and test safer cars in less time. The pharmaceutical industry is devel-

oping safe and effective drugs to combat cancer, AIDS, and other deadly diseases. Forecasters are better able to predict the weather and the paths of potentially dangerous storms.

Government-supported research is also responsible for many of the technologies that users of the information super highway are beginning to take for granted. the graphical mosaic interface that commercial enterprises are applying to make the Internet easier for consumers, the massively parallel computers that are being used to control massive amounts of network traffic and are critical components for the exploding world wide web, the very high speed computer chips that make searching very large distributed databases possible.

In conclusion, while the private sector has the primary responsibility for developing and deploying the NII and is investing significant amounts in R&D to advance the necessary technologies, there is a real need for a Government-industry partnership in HPCC and NII. There are certain pre-commercial NII research challenges, complex, long-term, and high risk, that are beyond the scope of any single company or even any industry sector.

We believe that the Federal Government with private sector input can provide a stimulus for this research and can promote the collaborative work among students, industries, academia, Government labs, and users. In some cases pilot projects and testbeds that apply and integrate advanced technologies must involve users and industry participants from different sectors such as computers, communications, software, database companies, as well as academic researchers.

We feel that this Federal support is not industrial policy. It creates people and ideas that succeed or fail in the marketplace over time. Federal support for pre-commercial HPCC research is essential to speed the application of information technologies in new areas, demonstrate the benefits they will make possible, and accelerate the development and adoption of standards by the private sector. The investments that the Federal Government makes today in research on high performance computing and networking technology will lay the foundation for an enhanced NII that will help solve problems and realize opportunities in tomorrow's homes, factories, universities, work places, classrooms, as well as boost the competitiveness of a wide range of U.S. industry sectors.

The United States has long enjoyed unparalleled success in the invention and application of information technology. The foundation for this incredible innovation engine has been the pioneering work done in our research universities and with Federal support. It would be a tragedy in this era of increased global competition if the United States cut back its support and funding of such critical research and training.

Thank you very much.

[The prepared statement of Dr. Baskett follows:]

Computer Systems Policy Project

TESTIMONY BY DR. FOREST BASKETT

CHAIRMAN, CSPP CHIEF TECHNOLOGISTS
SENIOR VICE PRESIDENT OF R&D and CHIEF TECHNOLOGY OFFICER,
SILICON GRAPHICS INC.

BEFORE THE SUBCOMMITTEE ON BASIC RESEARCH
U.S. HOUSE OF REPRESENTATIVES

TUESDAY, OCTOBER 31, 1995

TESTIMONY BY DR FOREST BASKETT
CHAIRMAN, CSPP CHIEF TECHNOLOGISTS
SENIOR VICE PRESIDENT OF R&D AND CHIEF TECHNOLOGY OFFICER,
SILICON GRAPHICS INC

BEFORE THE SUBCOMMITTEE ON BASIC RESEARCH
U S HOUSE OF REPRESENTATIVES

Good morning. My name is Forest Baskett. I am the Senior VP of Research and Development for Silicon Graphics and the Chairman of the Chief Technologists Committee of the Computer Systems Policy Project (CSPP). CSPP is a group of chief executive officers from the thirteen leading American computer systems companies. On their behalf as well as my own, I am truly pleased to be here today in support of the High Performance Computing and Communications (HPCC) program.

If there is anything I can leave you with today I would like it to be the following three points:

- *The HPCC program fosters U.S. innovation.* HPCC funding is driving innovation in the information technology industry and that innovation is happening in the United States and is leading the world. Without HPCC funding innovation will continue (at a slower rate) but it will be world wide. The playing field will become level.
- *If the HPCC program is severely cut, it is unlikely that the U.S. Information Technology Industry will fill the void.* The competitive nature of the current information technology industry means that companies that continue to survive do not and will not fund significant amounts of long range research. The long range nature of HPCC funding is why we are leading the world and we see no substitute for it.
- *Products from the information technology industry are pervasive.* Key U.S. industries benefit from information technology innovations. When our products become 100 times more powerful every seven years we change almost every other industry - manufacturing industries, service industries, financial industries, biotechnology industries, everything.

Silicon Graphics today is a multibillion dollar company that is part of the U.S. Information Technology industry. Some estimate that the information technology industry represents about 16% of our Gross Domestic Product and a much larger percentage of our total dollar volume of exports to the rest of the world. Computer systems companies like the members of the CSPP are not the whole of the information technology industry but computer systems are a key part of communications systems, a key part of information services, and certainly a key part of the software industry. As such we are keenly interested in the High Performance Computing and Communications program.

The History of SGI: The Role of Federal Investment

Silicon Graphics itself traces its existence back to federally funded science and technology. An ARPA program in Very Large Scale Integrated circuits had a major component at Stanford University. As part of that program Professor Jim Clark realized that it would be possible to do, on a single silicon chip, what had previously only been possible by a room full of traditional electronics, process and manipulate three-dimensional models in real time. The ARPA program made it possible to demonstrate that realization. That demonstration was just one of many that the ARPA program made possible and it was just one of many that then convinced the traditional capital markets that a real technological innovation had been achieved and that a new industry was feasible, based on that innovation. So Silicon Graphics was founded by Professor Clark and several of his graduate students as a result of the ARPA funded research they performed at Stanford. Many other companies were founded and several are as successful as Silicon Graphics. Since that time we at Silicon Graphics have continued to work with Stanford and other universities and we have continued to find new innovations that have continued to fuel our own growth and our leadership in world markets.

In fact, the members of CSPP believe that federal support of science and technology has been the innovation engine behind our industry. The discoveries that have been made, and that are continuing to be made, that have kept us so far ahead of the world in this industry are most often the result of work initiated as part of federal research programs such as the ARPA program I mentioned. Today many of those research programs in information technology are part of the HPCC program.

On the other hand, we recognize that the days of continued total growth in that support are over. Today's budget realities mean that we must be more thoughtful and more careful in our choices of what programs deserve support. So today I would like to describe some of the principles that we believe should help guide those choices.

The Need for Research Investment

Much of the innovation in our industry is made possible by the continued improvement in silicon device fabrication capabilities. As we look into the future we are confident in predicting that the 60% per year improvement in silicon device densities that we have been seeing for the last 30 years will continue for at least another 10 years. The 50% per year improvement in silicon device speeds will continue for at least another 10 years. The power of new innovations derives from this power base which is the product of device speeds and device densities. So we can easily expect to be able to do 100 times as much with a silicon device in seven years as we are able to do today. With this kind of change in the power of our tools, our accomplishments are only limited by our imagination and our opportunities for being imaginative.

The HPCC program provides many of those opportunities. Without those opportunities we spend most of our time simply trying to keep up with the technology in existing systems. Existing systems do get faster. Existing systems do get less expensive. But existing systems do not create whole new subindustries and industries. For example, will we be able to change the growth prospects for the commercial airline industry by inventing technologies that can make virtual meetings as good as or maybe even better than real meetings? It's possible. Could we invent a system that could schedule taxis for a city and its suburbs that could work so well and so efficiently and be so integrated with

mass transit that people would prefer it to private autos for most daily business travel? Would such a system change the nature of our cities? This problem is only about 1000 times more difficult than airline scheduling so it is possible I could try to go on except for my own limited imagination. The only other industry that even comes close to the information technology industry in its fundamental ability to innovate is the biotechnology industry.

Our industry today is a thriving and healthy one. It is vigorous and exciting. The US part of it is clearly leading the world. But it is also a competitive industry. That competition is part of why it is so vigorous. That competition has expanded our industry by bringing more new innovations to market more quickly and more cheaply. On the other hand, while that competition has expanded our industry it has also reduced our internal spending on R&D and caused most of that internal spending to be on development, not research. Many years ago when there was one dominant computer company and one dominant communications company, those companies had research enterprises that were the envy of the world. Those research enterprises existed for at least two reasons. One was that, partly because of the dominance of those companies, there was not a sufficient outside research base to draw on to fuel new development. The other was that those companies' dominant position meant that they could afford to support those research enterprises without endangering the existence of the company.

Today a big internal long range research program by a company is a burden that can easily kill the company in the competitive market place before any benefits of that research can be realized. We have become increasingly dependent on both university research and university trained professionals to fuel our continued growth and our continued leadership and our continued prosperity. The large private research enterprises of the past have been significantly scaled back in size and redirected at short term goals. Federal support of research is not what has caused the reduction in internal R&D and it's focus on development, competition has. Federal support of research has simply allowed us to continue to lead the world in this world wide competitive environment. If federal support were to disappear the playing field would become much more level. The rate of improvements in existing technologies and the rate of introduction of new innovations and new technologies and new industries would decrease. US leadership would be in danger.

Our companies, for the most part, are not looking to the federal government to support our internal R&D programs. The federal grants and contracts process is too slow for our fast moving industry. The federal reporting process for those grants and contracts is burdensome. The traditional capital markets are fast and efficient. Debt capital, equity capital, and venture capital are all available for worthy short term projects. It is the long term projects that are in need if we are to maintain our world leadership.

A significant issue for computer system companies and, I believe, for all information technology companies, is the current discussion of basic and applied research. We have difficulty with some of this discussion because of our inability to understand any reasonable distinction between the two. We are most sympathetic with Nobel laureate Lord Porter who said, "There are two kinds of research -- applied research and not-yet-applied research." We do understand the distinction between research and development. But we worry that an artificial distinction between basic and applied research will

endanger the research programs at our world class system of schools of engineering. The research programs that these schools of engineering have been a major source of innovation for our industry and certainly a major source of fresh, new innovative minds trained in the disciplines we most need.

The Benefits of the HPCC Program

There are several features of the HPCC program that we wholeheartedly support. Let me list them briefly but encourage you to consider them thoughtfully.

- Merit review for federal support
- Support for programs, not organizations or institutes or institutions
- Support that generates both research and training
- Multi-agency coordination and collaboration. Here the HPCC program is superb!
- Support for both science and engineering research, especially at universities
- Support for both theory and experiment. The experiments test the theories and train the students. The experiments nearly guarantee that technology transfer happens for the real innovations.

Significantly reducing the HPCC program and other National Information Infrastructure (NII) related research would seriously jeopardize the progress that this nation's researchers have been making in the critical areas of computing, software, and communications. The HPCC Program has been instrumental in expanding our knowledge about how to use high performance computers and high-speed networking to solve a range of scientific and engineering problems. It is also providing the research foundation for the National Information Infrastructure (NII), which will bring the benefits of information technologies to all Americans through applications in education, health care, access to government information and services, and digital libraries. University researchers are working to advance many of the technologies that are critical underpinnings of the NII, such as storage technologies, technologies to protect information security, ensure privacy and confidentiality, protect intellectual property, and authenticate information sources.

Many industrial sectors are using the technologies advanced through the HPCC Program to improve their manufacturing efficiency, shorten production times, reduce costs, and develop products that improve the quality of life:

- the auto industry is using high performance computers to design and test safer cars in less time;
- the pharmaceutical industry is developing safe and effective drugs to combat cancer, AIDS, and other deadly diseases; and
- forecasters are better able to predict the paths of potentially dangerous storms

Government-supported research is also responsible for many of the technologies that users of the information superhighway are beginning to take for granted:

- the graphic Mosaic interface that commercial enterprises are applying to make Internet use easier for consumers,
- the massively parallel computers that are being used to control massive amounts of network traffic and are critical components of trials for video on demand, and

- the very high speed computer chips that make searching very large, distributed databases possible

Conclusion

While the private sector has the primary responsibility for developing and deploying the NII and is investing significant amounts in R&D to advance the necessary technologies, there is a real need for a government-industry partnership in HPCC and NII. There are certain pre-commercial NII research challenges -- complex, long-term and high risk -- that are beyond the scope of any single company or industry sector.

We believe that the federal government, with private sector input, can provide a stimulus for this research and can promote collaborative work among industries, academia, government labs, and users. In some cases, pilot projects and testbeds that apply and integrate advanced technologies must involve users and industry participants from different sectors, such as computer, communications, software, and database companies, as well as academic researchers.

This federal support is not "industrial policy" -- it creates people and ideas that succeed or fail in the marketplace over time. Federal support for precommercial HPCC research is essential to speed the application of information technologies in new areas, demonstrate the benefits they will make possible, and accelerate the development and adoption of standards by the private sector. The investments that the federal government makes today in research on high performance computing and networking technology will lay the foundation for an enhanced NII that will help to solve problems and realize opportunities in tomorrow's homes, factories, universities, workplaces and classrooms, as well as boost the competitiveness of a wide range of U.S. industry sectors.

The United States has enjoyed unparalleled success in the invention and application of information technology. The foundation for this incredible innovation engine has been the pioneering work done in our research universities. It would be a tragedy, in this era of increased global competition, if the United States cut back its support and funding of such critical research and training.

Mr. SCHIFF. Thank you, Dr. Baskett. I assume all of the witnesses heard the bells and whistles going off behind you there. I am afraid they run our lives and we have a vote on the House floor.

So, I presume that the four of you would be available if we want to send written questions following this. I am not sure we will. I think all of your testimonies were very comprehensive.

I am just going to make one quick observation and that is, once again, as I said before to the last panel, I am a believer in Government-industry cooperation, but it still disturbs me when I hear industry just will not do this, industry will not invest in long-range R&D. Maybe we need to study with industry what could promote them to do that because the rather simple, "we are going to let the taxpayers do it for us" rankles me at a certain point even as a believer.

I am sorry but just because of the vote, I am going to have to recognize Mr. Geren.

Mr. GEREN. Thank you, Mr. Chairman. I would like to just pose a question and ask that you do answer it in writing, and this is a little bit outside of each of your areas of specialization, but an issue that we are going to be asked to confront is to reauthorize or not to reauthorize. Each of you has had experience with this type of research prior to the creation of HPCC in 1991 and since that time. There are supporters of the program who think we should not authorize because they think that perhaps that could bring mischief and some micromanagement by the Congress or perhaps even cuts that would be unfortunate. There are supporters of the program that think we got to have reauthorization if the program is going to stay as focused as we need it to be and that 12 different agencies could all go their separate ways and we would lose this coordination that some think has been very beneficial since 1991.

So, if you could draw on your experience—and it is not your technical experience so much, but your practical experience in seeing how this program works—and advise us whether in your opinion, as outsiders who have worked closely with it, whether or not you think reauthorization and the structure that that could lend to the program would be beneficial, or would it be better that we go back to the pre-1991 where the program has less structure?

There is not time for you to answer it now, but that is a very important question that we as a Committee are going to have to address. If you all could help us think through that, I would find that very helpful. Thank you, Mr. Chairman.

Mr. SCHIFF. Again, if you could get that to us in writing, say, in the next 10 days, I would appreciate that. Without objection, the record will stay open for 10 days.

I would like to thank this panel and the previous panel for your testimony, for your thoughts and your ideas and your expertise, and want to say that I think this has been a very instructive hearing. It is again a matter where we are all in agreement, but what is the best way to proceed. We appreciate your contribution to that.

The hearing is now adjourned.

[Whereupon, at 12:47 p.m., the Subcommittee was adjourned.]

[The following material was received for the record:]

Schlumberger

SCHLUMBERGER LIMITED
8311 NORTH RR 620
P.O. BOX 20015
AUSTIN, TEXAS 78720-0015
(512) 331-3713
FAX (512) 331-3760

JOHN D INGRAM
SCHLUMBERGER RESEARCH FELLOW

October 25, 1995

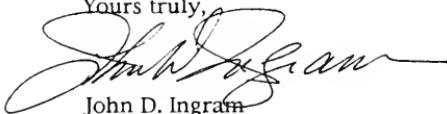
Mr. Steven H. Schiff
Chairman
Subcommittee on Basic Research
U. S. House of Representatives
Suite 2320 Rayburn House Office Building
Washington, DC 20515-6301

Dear Mr. Schiff;

In response to Representative Geren's question concerning the reauthorization of the HPCC Program I would offer the following comments.

First of all I do not feel qualified to comment on the legal and political aspects of reauthorization. I would, however, like to reiterate the importance of the research that HPCC represents. We are at a stage in which High Performance Computing can become a critical building block for the future of not only basic and applied research but for product development as well. There is much work to be done on the software environment to make this possible. It is research that spans agency lines and is the key to success in this fundamental area. The computer industry cannot be counted on for this work as it requires an effort that must provide standards and be independent of proprietary concerns. Whatever decision you reach it should contain the elements to assure consistent and rapid development of High Performance Computing Environments.

Yours truly,



John D. Ingram

wl.admin.on/Subcommittee

Princeton University

Office of the Provost

3 Nassau Hall

Princeton, New Jersey 08544-5264

November 28, 1995

Steven H. Schiff
Subcommittee on Basic Research
U.S. House of Representatives
Committee on Science
Suite 2320 Rayburn House Office Building
Washington, DC 20515-6301

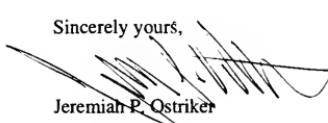
Dear Mr. Schiff:

I received your letter of November 13, requesting my response to Representative Geren's question, which he posed during the HPCC hearing on October 31, 1995, and which, I understand, will be included in the official hearing record. My response is as follows:

Congressman Geren has asked me whether the HPCC program should be re-authorized on the basis of my own practical experience before and after this program came into being. In one sense, this is an easy question. The program has been an enormous success! With little in the way of net additional resources, the HPCC initiative has allowed us in the scientific community to *focus* our efforts in a far better way on the major problems, and on the problems where real progress can be made. The inter-agency cooperation, while imperfect, has greatly facilitated such a coordinated approach. But, what is the best tactical approach to maintain the momentum of this program? Here, I have little to add which would be useful to the discussion. The United States governmental budgetary process is very difficult for an outsider to understand, and more so this year than ever. Furthermore, much of the resources devoted to this program was not from new funds but from the desire of the different agencies to reprogram some funds at hand to this exciting initiative, so new budgetary authority may not be necessary. Given this welter of complex and technical administrative issues, let me leave it to wiser and more experienced heads to determine how, operationally, one should act to maintain the very considerable strength we have achieved in high performance computing and communications, through the HPCC initiative.

If you should have any further questions, please do not hesitate to contact me. I can be reached at telephone #609-258-3026.

Sincerely yours,


Jeremiah P. Ostriker

2550 Garcia Avenue, MS UMTV29-Lobby
Mountain View, CA 94043-1100
415 336-2600

 **Sun Microsystems
Laboratories, Inc.**

A Sun Microsystems, Inc. Business

November 6, 1995

The Honorable Steven Schiff
Subcommittee on Basic Research
Committee on Science
U.S. House of Representatives
B374 Rayburn House Office Building
Washington, DC 20515
By Fax: (202) 225-7815

Dear Mr. Schiff:

On October 31, 1995, I testified before your committee at its hearing on High Performance Computing and Communications. I wish to add three responses to the record of that hearing. The first and second are personal responses to the question from Congressman Baker about how the government might help by getting out of the way. The third response is to questions by Congressmen Geren and Walker on the issue of basic vs. applied research. I offer this third response in my capacity as co-chair of the National Research Council committee on HPCCI.

Mr. Baker asked the witnesses as a group what the government might do to get out of the way of private industry to foster innovation. I have two answers to that question, one relating to industry and the other to the execution of the government research program itself.

My first answer relates to export controls for computing and communications equipment. For this answer I speak as a private citizen and as Vice President and Fellow at Sun Microsystems, a leading supplier of computer workstations and networks.

- 1) The government should get out of the way of private industry by eliminating export controls on computing and communications equipment.

The rapid pace of development in the computing and communications industry soon makes any capability-based export control regulation seem foolish. The super computer of a few years ago, considered then to be essential to designing military devices, is today available for home use and is freely exportable. Indeed, the government regularly revises export regulations to permit export of ever better equipment. Because of the rapid pace of information technology, export controls serve mainly to hamper U.S. industry in international competition.

If export controls for computing equipment must remain, base them on system cost or rarity rather than on system capability. It may be proper to control large and expensive systems, or rare systems of which only a few exist anywhere in the world. The government should get export controls out of the way of equipment makers in volume markets where the economic competition is most keen. Export control is impractical in volume markets and merely adds unnecessary cost to U.S. products.

My second answer relates to the government procurement process for supporting computing and communications research. For this answer I speak as a private citizen, as a former ARPA official involved in procuring research, and as an observer of many government supported research projects.

2) The government should get out of the way of research progress by streamlining its procurement process.

2a) The time from research proposal to grant should be short, a few months at most. I observe over and over again a waste of critical research talent devoted to working an inappropriate and cumbersome procurement process. The very essence of research is that its results are not known in advance, yet we burden our research procurement process with requirements more appropriate to purchasing material things than to supporting the development of new ideas. The lethargy of today's government research procurement process hampers the development of the fresh ideas our nation needs to remain competitive.

2b) Encourage top researchers to serve government in the mission agencies. Our government needs to borrow leading research talent from industry and academia, and should expect such people to return to private life after serving for a few years in government. Such a flow of people from research projects in academia and industry to research management in government would serve our nation well.

Today a burden of regulation hampers such people not only while they serve in government but also after they leave to return to private life. To compete effectively for top talent, government must make its positions as attractive as possible. Getting conflict of interest provisions out of the way and providing access to support of continuing research after government service would encourage more top research people to serve in government. Consider the plight of a researcher who devotes two or three years to government service. It may take two or more years afterwards to rebuild his or her research program, and he or she may even be denied access to government research support because of recent government service. Top people prefer private industry or academia to government service, but I believe that this preference reflects excessive regulation rather than salary.

2c) Sole source procurement often makes sense for research. Individuals discover things because of their unique capability and experience. Our present research procurement process ignores this very important fact. Good ideas presented by good people as unsolicited proposals for government research support should be considered on the merits of the ideas and the people. Capable government officials in the mission agencies can and should decide

which people and which ideas to support. We should trust our officials to choose wisely by giving them sole source authority for research work. We should hold them accountable for the collective fruits of their research programs. We must recognize that not all ideas will reach fruition, but a well run research program in a mission agency will have a better than average overall track record.

My third response relates to Mr. Geren's question about recommendation 10 from the National Research Council report, Evolving the HPCCI to Support the Nation's Information Infrastructure. Recommendation 10 of our report concerns the "grand challenge" component of the HPCCI. I respond as co-chair of the committee responsible for that report.

3) We should fine tune the decision process for "grand challenge" projects. During the early period of the High Performance Computing and Communications Initiative several interdisciplinary teams formed to explore the uses of new technology to scientific problems. Such teams served two purposes. First, they could explore the new computing capability in several different scientific contexts, thus contributing directly to the HPCCI goals. Second, they could use the new computing capability to do good science, thus contributing to the advancement of scientific fields other than computing. Both were worthy goals.

Our committee continues to support both goals as deserving of federal support. Many of the "grand challenge" projects have enjoyed a measure of success towards both goals; they have done good science and learned much about using the new parallel form of computer. Indeed, most of the computing power being used today in the NSF super computer centers comes from the newer form of parallel computers, see Figure E.1, page 109 of our report.

Our recommendation 10 focuses narrowly on the use of HPCCI funds rather than broadly on government support as a whole. We recommend that HPCCI funds, which should properly focus on computing and communications, be used for projects that explore the frontiers of computing and communications as well as contributing to other scientific frontiers. We applaud support of projects that have worthy scientific goals regardless of their contribution to high performance computing and communications, but we believe that support for projects should be included in HPCCI only if they contribute to HPCCI's goals.

With regard to Mr. Walker's question about support for basic research I wish to reiterate that our NRC report's number one recommendation is to continue support of a basic research program in computing and communications. Such support is vital to the continuing health of the U. S. information industry. Moreover, our report points out the important links between our information industry and many other aspects of our economy and our society including manufacturing, health care, education, research, electronic commerce, financial services, entertainment, information retrieval, and quality of life.

Support of basic research is a continuing activity that should remain stable regardless of any initiative. There has been a natural tendency in the mission agencies to give many activities an HPCCI label. Thus, for example, nearly all of the NSF program in computing and communications bears the HPCCI label. Our committee saw as dangerous the trend to use HPCCI as an umbrella because many basic activities must continue beyond the boundary of

298

3 9999 05984 233 4

any initiative. We believe that the broader question of how best to support basic research is confused by the existence of the HPCCI initiative. Recommendation 12 of our report seeks to distinguish programs that should be on-going from those that are uniquely a part of the HPCC initiative itself.

I thank you again for affording me the opportunity to testify before your subcommittee. I am always happy to discuss the benefits of today's federal support for basic research because I believe that today's research finds the new ideas vital to tomorrow's economic well-being for us and our children. We live in a highly competitive world. Those of us in private industry depend on federal research, largely in academia, for fresh new ideas. I hope that you will see fit to keep our nation's storehouse of ideas well stocked. The number one recommendation of our NRC group is to keep our intellectual seed corn alive.

Sincerely,



Ivan E. Sutherland

enclosure: (1) Figure E.1, p.109, *Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure*.

SDS 95:0390

Supercomputer Usage at NSF Centers

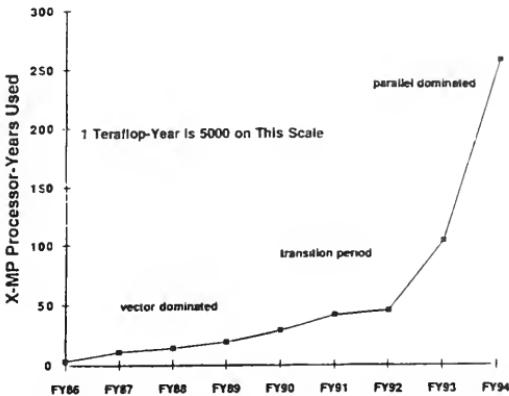


Table E.1 shows the growth in the number of users and in the availability of cycles at the NSF supercomputer centers from 1986 to 1994. See also Figure E.1. The increase in capacity in 1993 was owing mainly to the introduction of new computing architectures. The slight decrease in the number of users reflects the centers' effort to encourage users able to meet their computational needs with the increasingly powerful workstations of the mid-1990s to use their own institutional resources.

FIGURE E.1 Total historical usage of all high-performance computers in the NSF MetaCenter. This graph shows the total annual usage of all high-performance computers in MetaCenter facilities. Particularly striking is the growth since 1992, when microprocessors in various parallel configurations began to be employed. All usage has been converted to equivalent processor-years for a Cray Research Y-MP, the type of supercomputer that the NSF centers first installed in 1985-1986.

TABLE E.1 Supercomputer Usage at National Science Foundation Supercomputer Centers, 1986 to 1994

Fiscal Year	Active Users	Usage in Normalized CPU Hours*
1986	1,358	29,485
1987	3,326	95,752
1988	5,069	121,615
1989	5,975	165,950
1990	7,364	250,628
1991	7,887	361,037
1992	8,578	398,932
1993	7,730	910,088
1994	7,431	2,249,562

*Data prior to May 1990 include the John von Neumann Center.

ISBN 0-16-052412-1



9 0 0 0 0

9 780160 524127